

# UTILITY PATENT APPLICATION TRANSMITTAL

Submit an original and a duplicate for fee processing  
(Only for new nonprovisional applications under 37 CFR 1.53(b))

ADDRESS TO:

Assistant Commissioner for Patents  
Box Patent Application  
Washington, D.C. 20231

Attorney Docket No. 99,226

First Named Inventor Michael S. Borella

Express Mail No. EL007910275US

Total Pages

## APPLICATION ELEMENTS

1. ☒ Transmittal Form in duplicate with Fee
2. ☒ Specification (including claims and abstract) [Total Pages 83]
3. ☒ Formal Drawings [Total Sheets 27]
4. ☐ Oath or Declaration [Total Pages ]
  - a. ☐ Newly executed
  - b. ☐ Copy from prior application

**[Note Boxes 5 and 17 below]**

  - i. ☐ Deletion of Inventor(s) Signed statement attached deleting inventor(s) named in the prior application
5. ☐ Incorporation by Reference: The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.
6. ☐ Microfiche Computer Program
7. ☐ Nucleotide and/or Amino Acid Sequence Submission
  - a. ☐ Computer Readable Copy
  - b. ☐ Paper Copy
  - c. ☐ Statement verifying above copies

## ACCOMPANYING APPLICATION PARTS

8. ☐ Assignment Papers
9. ☐ Power of Attorney
10. ☐ English Translation Document (if applicable)
11. ☐ Information Disclosure Statement (IDS)
  - ☐ PTO-1449 Form
  - ☐ Copies of IDS Citations
12. ☐ Preliminary Amendment
13. ☒ Return Receipt Postcard (Should be specifically itemized)
14. ☐ Small Entity Statement(s)
  - ☐ Enclosed
  - ☐ Statement filed in prior application; status still proper and desired
15. ☐ Certified Copy of Priority Document(s)
16. ☒ Other: Certificate of Express Mail

17. **If a CONTINUING APPLICATION**, check appropriate box and supply the requisite information:  
☐ Continuation ☐ Divisional ☐ Continuation-in-part of prior application Serial No.

## APPLICATION FEES

BASIC FEE				\$760.00
CLAIMS	NUMBER FILED	NUMBER EXTRA	RATE	
Total Claims	20	-20=	x \$18.00	\$
Independent Claims	3	- 3=	x \$78.00	\$
<input type="checkbox"/> Multiple Dependent Claims(s) if applicable			+\$270.00	\$
Total of above calculations =				\$760.00
Reduction by 50% for filing by small entity =				\$( )
<input type="checkbox"/> Assignment fee if applicable			+ \$40.00	\$
TOTAL =				\$760.00

**UTILITY PATENT APPLICATION TRANSMITTAL**


Attorney Docket No. 99,226

18. ☐ Please charge my Deposit Account No. 13-2490 in the amount of \$ .
19. ☒ A check in the amount of \$760.00 is enclosed.
20. The Commissioner is hereby authorized to credit overpayments or charge any additional fees of the following types to Deposit Account No. 13-2490:
- a. ☒ Fees required under 37 CFR 1.16.
  - b. ☒ Fees required under 37 CFR 1.17.
  - c. ☒ Fees required under 37 CFR 1.18.
21. ☒ The Commissioner is hereby generally authorized under 37 CFR 1.136(a)(3) to treat any future reply in this or any related application filed pursuant to 37 CFR 1.53 requiring an extension of time as incorporating a request therefor, and the Commissioner is hereby specifically authorized to charge Deposit Account No. 13-2490 for any fee that may be due in connection with such a request for an extension of time.

**22. CORRESPONDENCE ADDRESS**

Name	McDonnell Boehnen Hulbert & Berghoff
Address	32 <sup>nd</sup> Floor, 300 South Wacker Drive
City, State, Zip	Chicago, Illinois 60606

**23. SIGNATURE OF APPLICANT, ATTORNEY, OR AGENT REQUIRED**

Name	Sean M. Sullivan Registration No. 40,191
Signature	
Date	August 27, 1999


UTILITY (Rev. 11/18/97)

**CERTIFICATE OF MAILING BY "EXPRESS MAIL"**  
(NEW PATENT APPLICATION)

Express Mail No. EL007910275US

Deposited August 27, 1999

I hereby certify that the attached correspondence, identified below, is being deposited with the United States Postal Service as "Express Mail Post Office to Addressee" under 37 CFR § 1.10 on the date indicated above and is addressed to the Asst. Commissioner for Patents, Box New Application, Washington, DC 20231.

By:   
(person actually depositing)

Patent Application of: Michael S. Borella, Gary Jaszewski, Danny M. Nessett

Title: Method and System for Controlling Attacks on Distributed Network Address Translation Enabled Networks

- (1) Utility Patent Application Transmittal in duplicate;
- (2) Specification including claims and Abstract (83 sheets);
- (3) Formal Drawings (27 sheets);
- (4) Filing Fee Check in the amount of \$760.00
- (5) Return postcard.

Attorney Docket No.: 99,226

**APPLICATION FOR A UNITED STATES PATENT**  
**UNITED STATES PATENT AND TRADEMARK OFFICE**  
(Case No. 99,226)

**Title: METHOD AND SYSTEM FOR CONTROLLING ATTACKS ON  
DISTRIBUTED NETWORK ADDRESS TRANSLATION ENABLED  
NETWORKS**

**Inventors:** Michael S. Borella, a citizen of the United States, and a resident of Naperville,  
Illinois;

Gary Jaszewski, a citizen of the United States, and a resident of Los Gatos,  
California; and

Danny M. Nessett, a citizen of the United States, and a resident of Fremont,  
California.

**Assignee:** 3Com Corporation  
5400 Bayfront Plaza  
Santa Clara, CA 95052

**Attorney:** Sean M. Sullivan, Reg. No. 40,191  
McDonnell, Boehnen, Hulbert & Berghoff  
300 South Wacker Drive  
Chicago, Illinois 60606  
Tel. No. (312) 913-0001

## CROSS REFERENCES TO RELATED APPLICATIONS

This application is a Continuation-In-Part of U.S. Application No. 09/035,600, filed on  
5 March 5, 1998, and U.S. Application No. 09/270,967, filed on March 17, 1999, both of which  
are specifically incorporated in their entirety herein by reference.

## FIELD OF INVENTION

10 This invention relates to computer networks. More specifically, it relates to a method and  
system for distributed network address translation with network security features that can be  
used to control and limit disruptions caused by Denial of Service attacks.

## BACKGROUND OF THE INVENTION

15 The Internet Protocol ("IP") is an addressing protocol designed to facilitate the routing of  
traffic within a network or between networks. The IP is used on many computer networks  
including the Internet, intranets and other networks. Current versions of IP such as Internet  
Protocol version-4 ("IPv4") are becoming obsolete because of limited address space. With a 32-  
bit address-field, it is possible to assign  $2^{32}$  different addresses, which is 4,294,967,296, or  
20 greater than 4 billion globally unique addresses.

However, with the explosive growth of the Internet and intranets, IP addresses using a  
32-bit address-field may soon be exhausted. Internet Protocol version-6 ("IPv6") proposes the  
use of a 128-bit address-field for IP addresses. However, a large number of legacy networks  
including a large number of Internet subnets will still be using older versions for IP with a 32-bit  
25 address space for many years to come.

Network Address Translation ("NAT") has been proposed to extend the lifetime of Internet Protocol version 4 and earlier versions of Internet Protocol by allowing subnets to exist behind a single or small number of globally unique IP addresses (see e.g., "The IP Network Address Translator", by P. Srisuresh and K. Egevang, Internet Engineering Task Force ("IETF"), Internet Draft <draft-rfced-info-srisuresh-05.txt>, February 1998). A single global IP address is used for communication with external networks such as the Internet. Internally, a sub-network ("subnet") uses local addressing. Local addressing may be either any addressing scheme that is different from IP addressing, or a non-unique usage of IP addresses. In either case, local addresses on a subnet are not used on the external, global Internet. When a device or node using local addressing desires to communicate with the external world, its local address is translated to a common external IP address used for communication with an external network by a NAT device. That is, NAT allows one or more global IP addresses to be shared among a larger number of local addresses.

There are several problems associated with using NAT to extend the life of the IP. NAT interferes with the end-to-end routing principal of the Internet that recommends that packets flow end-to-end between network devices without changing the contents of any packet along a transmission route (see e.g., "Routing in the Internet," by C. Huitema, Prentice Hall, 1995, ISBN 0-131-321-927).

Current versions of NAT replace a local network address in a data packet header with an external global network address on outbound traffic, and replace an external network address in a data packet header with a local network address on inbound traffic. This type of address translation is computationally expensive, causes security problems by preventing certain types of

encryption from being used, or breaks a number of existing applications in a network that cannot provide NAT (e.g., File Transfer Protocol ("FTP")).

Current versions of NAT may not gracefully scale beyond a small subnet containing a few dozen nodes or devices because of the computational and other resources required. NAT potentially requires support for many different internal network protocols be specifically programmed into a translation mechanism for external protocols in a NAT device such as a NAT router.

Computational burdens placed on a NAT router may be significant and degrade network performance, especially if several NAT-enabled sub-networks share the same NAT router. In a worst case scenario, a NAT router translates every inbound and outbound data packet. When NAT is used to translate a TCP/IP or UDP/IP data packet, the packet's IP, TCP or UDP checksums are recalculated.

As is known in the art, TCP ("TCP") and UDP are often used over IP in computer networks. TCP provides a connection-oriented, end-to-end reliable protocol designed to fit into a layered hierarchy of protocols that support multi-network applications. UDP provides a transaction oriented datagram protocol, where delivery and duplicate packet protection are not guaranteed.

When a port in a TCP or UDP header is translated, the packet's TCP or UDP checksums are also recalculated. This further increases the computational cost of translation in a NAT router.

When an IP address or port is translated with NAT, a new length may result for the data packet and a possible change in a TCP sequence number. A running sequence number offset

(i.e., a delta) must then be maintained throughout the remainder of the connection. This delta must be applied to future traffic, including acknowledgment numbers further increasing computational time in a NAT router.

In addition to TCP or UDP, a NAT router may also translate network addresses, ports, change lengths and maintain sequence numbers for a number of different protocols that may use an IP address or port number (e.g., FTP, H.323, H.324, CUSeeME, RealAudio, Internet Relay Chat and others). This translation may further increase computational time in a NAT router.

The IP is used on global computer networks such as the Internet, and on many private networks such as intranets and Virtual Private Networks. It is often desirable to protect information sent with the IP using different types of security. Using security with the IP allows private or sensitive information to be sent over a public network with some degree of confidence that the private or sensitive information will not be intercepted, examined or altered.

IPSEC is a protocol for implementing security for communications on networks using the IP through the use of cryptographic key management procedures and protocols.

Communications between two endpoints of an IP traffic flow are made end-to-end-secure by the IPSEC protocol on an individual IP packet-to-packet basis. IPSEC protocol entities at connection endpoints have access to, and participate in, critical and sensitive operations that make a common connection secure.

IPSEC currently includes two security services, each having an associated header that is added to an IP packet that is being protected. The two security services include an Authentication Header ("AH") and an Encapsulating Security Payload ("ESP") header. The Authentication Header provides authentication and integrity protection for an IP packet. The



Encapsulating Security Payload header provides encryption protection and authentication for an IP packet.

The IPSEC protocol headers are identified in a protocol field of an IP data packet header.

The IPSEC protocol header specifies the type (e.g., Authentication Header or Encapsulating

Security Payload) and contains a numerical value called the Security Parameter Index ("SPI").

The Security Parameter Index together with a destination IP address and Internet Security protocol form a unique identifier used by a receiving system to associate a data packet with a construct called a "security association." The Security Parameter Index is used by the receiving system to help correctly process an IP packet (e.g., to decrypt it, or to verify its integrity and authenticity).

IPSEC establishes and uses a Security Association ("SA") to identify a secure channel between two endpoints. A Security Association is a unidirectional session between two termination endpoints. Two termination endpoints of a single Security Association define a logical session that is protected by IPSEC services. One endpoint sends IP packets, and a second endpoint receives the IP packets. Since a Security Association is unidirectional, a minimum of two Security Associations is required for secure, bi-directional communications. It is also possible to configure multiple layers of IPSEC protocols between two endpoints by combining multiple Security Associations.

There are several problems associated with using current versions of NAT when security is required and the IPSEC protocol is used. Current versions of NAT violate certain specific principles of the IPSEC protocol that allow establishment and maintenance of secure end-to-end connections of an IP network.

A NAT router typically needs to modify an IP packet (e.g., network ports, etc.).

However, once an IP packet is protected by IPSEC, it must not be modified anywhere along a path from an IPSEC source to an IPSEC destination. Most NAT routers violate IPSEC by modifying, or attempting to modify individual IP packets.

5 Even if a NAT router does not modify data packets it forwards, it must be able to read network port numbers (e.g., TCP, UDP, etc.) in the data packets. If certain IPSEC features are used (e.g., Encapsulated Security Payload ("ESP")), the network port numbers are encrypted, so the NAT router typically will not be able to use the network ports for NAT mapping.

10 Local host network devices on a Local Area Network ("LAN") that use NAT typically possess only local, non-unique IP addresses. The local non-unique IP addresses do not comprise a name space that is suitable for binding an encryption key (e.g., a public key) to a unique entity. Without this unique binding, it is not possible to provide necessary authentication for establishment of Security Associations. Without authentication, an endpoint of a connection cannot be certain of the identity of another endpoint, and thus cannot establish a secure and  
15 trusted connection.

Local host network devices on the LAN that use NAT may also be susceptible to denial of service attacks from external network devices that have not established SAs with the local host network devices. By the external network devices transmitting data packets using an SPI that belongs to the local host network devices, as well as an IP address that belongs to the NAT  
20 router and is shared with the local host network devices, these packets will be forwarded by the NAT router to the local host network devices. While the local host network devices may discard these packets upon receipt, the external network devices may transmit hundreds or thousands of

packets in rapid succession, thereby swamping resources of the LAN, its local host network devices, and/or the NAT router. This swamping of resources is well-known in the art as a Denial of Service (DoS) attack. DoS attacks can cause disruptions, or even complete breakdowns, in communications among local host network devices, and between local host network devices and external network devices.

Thus, it is desirable to provide a method for DNAT with IPSEC that can control and limit disruptions caused by DoS attacks. As with ordinary DNAT, this method should not increase the burden on a router or other network device that provides the address translation. In addition, this method should also allow IPSEC to be used with DNAT to provide secure communications between internal and external network devices.

## SUMMARY OF THE INVENTION

The present invention provides a method for distributed network address translation with security comprising the steps of providing a first network device and a second network device on a first network, and establishing a security association between the first network device and a third network device on a second network external to the first network. The method of the present invention also comprises the step of specifying an external address of the third network device for the security association. The method of the present invention further comprises the steps of storing the external address in a table on the second network device, and mapping at least one of an internal address and a security value to the external address in the table.

In addition, the present invention provides a method for distributed network address translation with security comprising the step of providing a first network device and a second network device on a first network, and a third network device on a second network external to the first network. The method of the present invention also comprises the steps of sending a packet having an external address and a security value from the third network device to the first network device, and intercepting the packet with the second network device. The method of the present invention further comprises the steps of determining whether the security value of the packet has been allocated to the first network device, and determining whether the external address of the packet has been specified by the first network device as being valid. Moreover, the method of the present invention comprises the step of sending the packet from the second network device to the first network device if the security value has been allocated to the first network device and the external address of the packet has been specified by the first network device as valid.

The present invention also provides a system for distributed network address translation with security comprising a routing network device that uses distributed network address translation with security to provide routing services for a plurality of internal and external network devices. The system of the present invention also comprises an established security association table associated with the routing network device for storing external addresses of external network devices that have established security associations with internal network devices. The established security association table may also be used for mapping external addresses that have been specified as valid by the internal network devices to one of internal network addresses and security values for established security associations.

The foregoing and other features and advantages of a preferred embodiment of the present invention will be more readily apparent from the following detailed description, which proceeds with references to the accompanying drawings.

## BRIEF DESCRIPTION OF THE DRAWINGS

Preferred embodiments of the present inventions are described with reference to the following drawings, wherein:

FIG. 1 is a block diagram illustrating a network system for distributed address translation;

FIG. 2 is a block diagram illustrating a protocol stack for a network device;

FIG. 3 is a block diagram illustrating a port allocation protocol ("PAP");

FIG. 4 is a block diagram illustrating a PAP request message layout;

FIG. 5 is a block diagram illustrating a PAP response message layout;

FIG. 6 is a block diagram illustrating a PAP invalidate message layout;

FIG. 4A is a block diagram illustrating a PAP security request message layout;

FIG. 5A is a block diagram illustrating a PAP security response message layout;

FIG. 6A is a block diagram illustrating a PAP security invalidate message layout;

FIG. 7 is a block diagram illustrating a PAP combined network address layout;

FIG. 8 is a block diagram illustrating a PAP port-to-internal network address table layout;

FIG. 9 is a flow diagram illustrating a method for allowing distributed NAT;

FIG. 10 is a flow diagram illustrating a method for distributed network address translation;

FIG. 11 illustrates a source port transition table layout;

FIG. 12 illustrates an IP address translation table layout;

FIG. 13 illustrates a method for outbound distributed network address translation using port translation;

FIG. 14 illustrates a method for inbound distributed network address translation using port translation;

FIG. 15 is a block diagram illustrating an IP packet header format;

FIG. 16 is a block diagram illustrating an IPSEC Authentication Header format;

5 FIG. 17 is a block diagram illustrating an Encapsulating Security Payload packet format;

FIG. 18 is a block diagram illustrating end-to-end security between two endpoints over an IP network;

FIG. 19 is a flow diagram illustrating a method for distributed network address translation with security;

10 FIG. 20 is a flow diagram illustrating a method for distributed network address translation with security;

FIG. 21 is a block diagram illustrating a SPI-to-internal network address table layout;

FIG. 22 is a flow diagram illustrating a method for providing a security association using distributed network address translation;

15 FIG. 23 is a flow diagram illustrating a method for distributed network address translation using security;

FIG. 24 is a flow diagram illustrating a method for distributed network address translation using security;

20 FIG. 25 is a flow diagram illustrating a method for distributed network address translation with security;

FIG. 26 is a block diagram illustrating a network system for distributed address translation with an additional external network device;

FIG. 27 is a flow diagram illustrating a method for controlling denial of service attacks using distributed network address translation with security;

FIGS. 28A-B are block diagrams illustrating a PAP external address validating message layout and a PAP external address invalidating message layout, respectively;

5        FIGS. 29A-C are block diagrams illustrating established security associations tables; and  
FIG. 30 is a flow diagram further illustrating the method of FIG. 27.



## DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

### **Exemplary network system**

FIG. 1 is a block diagram illustrating an exemplary network system 10 for one preferred  
embodiment of the present invention. The network system 10 includes a first computer network  
12 with multiple network devices (14, 16, 18, 20, 22, 24) and a router 26 to route data packets to  
another external computer network. The multiple network devices include any of computers (14,  
18), printers 16, facsimile devices 24, hand-held devices 20, telephones 22 or other network  
devices not illustrated in FIG. 1. The first computer network 12 has an external common  
network address 28 (e.g., a global IP address 198.10.20.30) to identify the first network 12 to an  
external computer network such as a second computer network 30 and/or a third computer  
network 32 external to first computer network 12. The multiple network devices (14, 16, 18, 20,  
22, 24, and 26) have an internal network address (i.e., a private network address) on the first  
computer network 12 (e.g., 10.0.0.x explained below). In one preferred embodiment of the  
present invention, a network access service provider 34 with a router 36 routes data packets  
to/from first computer network 12 to second computer network 30 and/or third computer  
network 32 through a second network switch 38 and/or a third network switch 40. In another  
embodiment of the present invention, the first computer network exemplary 12 is connected  
directly to second computer network 30.

In one preferred embodiment of the present invention, the first computer network 12 is a  
Small Office/Home Office (“SOHO”) Local Area Network (“LAN”), also called a “legacy”  
LAN. First computer network 12 is also called a “stub” network. As is known in the art, a stub

network typically includes multiple network devices using a common external network address to communicate with an external network such as the Internet. The second network 30 is the Internet or an intranet, and the third network 32 is a Public Switched Telephone Network ("PSTN"). However, other network types and network components can also be used and the present invention is not limited to the network types and network components described for this preferred embodiment. The present invention can be used with virtually any network using the IP or other protocols in the IP suite.

Network devices and routers for preferred embodiments of the present invention include network devices that can interact with network system 10 based on standards proposed by the Institute of Electrical and Electronic Engineers ("IEEE"), International Telecommunications Union-Telecommunication Standardization Sector ("ITU"), Internet Engineering Task Force ("IETF"), or Wireless Application Protocol ("WAP") Forum. However, network devices based on other standards could also be used. IEEE standards can be found on the World Wide Web at the Universal Resource Locator ("URL") "www.ieee.org." The ITU, (formerly known as the CCITT) standards can be found at the URL "www.itu.ch." IETF standards can be found at the URL "www.ietf.org." The WAP standards can be found at the URL "www.wapforum.org."

An operating environment for network devices and routers of the present invention include a processing system with at least one high speed Central Processing Unit ("CPU") and a memory. In accordance with the practices of persons skilled in the art of computer programming, the present invention is described below with reference to acts and symbolic representations of operations or instructions that are performed by the processing system, unless

indicated otherwise. Such acts and operations or instructions are referred to as being "computer-executed" or "CPU executed."

It will be appreciated that acts and symbolically represented operations or instructions include the manipulation of electrical signals or biological signals by the CPU. An electrical system or biological system represents data bits which cause a resulting transformation or reduction of the electrical signals or biological signals, and the maintenance of data bits at memory locations in a memory system to thereby reconfigure or otherwise alter the CPU's operation, as well as other processing of signals. The memory locations where data bits are maintained are physical locations that have particular electrical, magnetic, optical, or organic properties corresponding to the data bits.

The data bits may also be maintained on a computer readable medium including magnetic disks, optical disks, organic memory, and any other volatile (e.g., Random Access Memory ("RAM")) or non-volatile (e.g., Read-Only Memory ("ROM")) mass storage system readable by the CPU. The computer readable medium includes cooperating or interconnected computer readable medium, which exist exclusively on the processing system or be distributed among multiple interconnected processing systems that may be local or remote to the processing system.

In NAT schemes known in the art, the router 26 translates an internal network address such as an internal network address used on the first computer network 12 to an external network address such as a network address for outgoing traffic to the second network 30 or the third network 32. The router 26 also translates an external network address to an internal network address for incoming traffic from the second network 30 or the third network 32. A NAT router assumes the entire computation burden for network address translation. For large subnets, the

NAT router becomes a bottleneck. In the worst case, every packet passing through the NAT router will require address translation. For more information on NAT for the IP see "The IP Network Address Translator (NAT)," Internet Engineering Task Force ("IETF") Request For Comments ("RFC") RFC-1631, "NAT Bypass for 'End 2 End' sensitive applications," by G. Tsirtsis and A. O'Niell, IETF Internet Draft, <draft-tsirtsis-nat-bypass-00.txt>, January 1998, or "The IP Network Address Translator", by P. Srisuresh and K. Egevang, Internet Engineering Task Force ("IETF"), Internet Draft <draft-rfcd-info-srisuresh-05.txt>, February 1998.

In one preferred embodiment of the present invention, Distributed Network Access Translation ("DNAT") is used. Network devices (14, 16, 18, 22 and 24) on the first computer network 12 request a set of locally unique ports from the router 26 for external communications with the external second network 30 or the third network 32. A locally unique port is unique inside of the first computer network 12 and typically is not unique outside of first computer network 12. Locally unique ports may be used for mobile network devices, such as device 20 using Mobile IP, that are not permanently attached to the first computer network 12. A mobile network device may physically relocate to another location and attach to a foreign computer network (i.e., other than home computer network 12).

The network devices (14, 16, 18, 20, 22, 24) replace default or ephemeral ports with the locally unique ports and use a combination network address including a locally unique port and a common external network address (e.g., an IP address) for communications with the external networks 30 and 32. A default port is typically statically assigned. An ephemeral port is typically dynamically assigned for a specified duration of time.

### **DNAT Protocol Stack**

FIG. 2 is a block diagram illustrating a layered protocol stack 42 for a network device from the first computer network 12 used for DNAT. The layered Protocol stack 42 is described with respect to IP suites comprising from lowest-to-highest, a link, network, transport and application layer. However, more or fewer layers could also be used, and different layer designations could also be used for the layers in the protocol stack 42 (e.g., layering based on the Open Systems Interconnection ("OSI") model).

The network devices (14, 16, 18, 20, 22, and 24) are connected to the first computer network 12 with Network Interface Card ("NIC") device drivers 44 for the hardware network devices connecting the network devices to the computer network 12. Above the network interface card device drivers 44 is a network layer 46 (also called the Internet Layer for IP suites). The network layer 46 includes an IP layer 48. As is known in the art, IP 48 is an addressing protocol designed to route traffic within a network or between networks. IP layer 48, hereinafter IP 48, is described RFC-791, incorporated herein by reference.

Above network layer 46 is a transport layer 50. The transport layer 50 includes a Port Allocation Protocol ("PAP") layer 52, an Internet Group Management Protocol ("IGMP") layer 54, a Control Message Protocol ("ICMP") layer 56, a TCP layer 58 and a UDP layer 60. However, more or fewer protocols could also be used.

The PAP layer 52 allocates locally unique ports to a network device. In one embodiment of the present invention, the PAP layer 52, is a separate protocol layer in the network layer 46. In another embodiment of the present invention, the PAP layer 52 is implemented as part of the ICMP layer 50 and is not a separate protocol layer. In yet another embodiment of the present invention, PAP layer 52 is run over either a TCP or UDP. PAP layer 52 is explained below.

IGMP layer 54, hereinafter IGMP 54, is responsible for multicasting. For more information on IGMP 54 see RFC-1112, incorporated herein by reference.

ICMP layer 56, hereinafter ICMP 56, is used for IP control. The main functions of ICMP 56 include error reporting, reachability testing (e.g., "pinging"), route-change notification, performance, subnet addressing and other maintenance. For more information on ICMP 56 see RFC-792, incorporated herein by reference.

TCP layer 58, hereinafter TCP 58, provides a connection-oriented, end-to-end reliable protocol designed to fit into a layered hierarchy of protocols which support multi-network applications. TCP 58 provides for reliable inter-process communication between pairs of processes in network devices attached to distinct but interconnected networks. For more information on TCP 58 see RFC-793, incorporated herein by reference.

UDP layer 60, hereinafter UDP 60, provides a connectionless mode of communications with datagrams in an interconnected set of computer networks. UDP 60 provides a transaction oriented datagram protocol, where delivery and duplicate packet protection are not guaranteed. For more information on UDP 60 see RFC-768, incorporated herein by reference. Both TCP 58 and UDP 60 are not required in protocol stack 42. Either TCP 58 or UDP 60 can be used without the other.

Above transport layer 56 is an application layer 62 where application programs to carry out desired functionality for a network device reside. For example, the application programs for the network device 16 may include printer application programs, while application programs for the network device 24 may include facsimile application programs more or fewer protocol layers can also be used in the protocol stack 42.

## DNAT Protocol

FIG. 3 is a block diagram illustrating a Port Allocation Protocol ("PAP") 64. PAP 64 is implemented in a separate PAP layer 52 or as an integral part of ICMP 50 in the protocol stack 42 (FIG. 2). PAP 64 includes a PAP request message 66, a PAP response message 68, a PAP invalidate message 70 and a combination network address 72. PAP 64 also includes a PAP security request message 67, a PAP security response message 69, a PAP security invalidate message 71. The PAP security messages 67, 69, 71 are used for IPSEC and are explained below. In one preferred embodiment of the present invention, fields in the PAP messages (66, 68, 70, 67, 69, 71) follow standard ICMP 50 message format. However, other message layouts (i.e., Non-ICMP 50 message format) and more or fewer messages could also be used for PAP 64 messages.

In one preferred embodiment of the present invention, the PAP request message 66 is sent from a network device (14, 16, 18, 20, 22, and 24) to the router 26, to request a block of locally unique port numbers. In another embodiment of the present invention, the PAP 64 is used with another network device (e.g., a port server or other network device separate from the router 26). In another preferred embodiment of the present invention, the PAP 64 is used to request a block of Security Parameter Indexes ("SPI") that will be used to establish Security Associations ("SA") when IPSEC is used. Use of the SPIs will be explained below.

FIG. 4 is a block diagram illustrating a PAP request message layout 74. A type-field 76 is one-byte and has a value (e.g., 32) for requesting locally unique ports. A code-field 78 is one-byte and has a value of zero for ports under 10,000 and a value of one for ports 10,000 or above. A checksum-field 80 is two-bytes, and has a value of a 1's complement sum of the entire PAP

request message 66 layout 74. As is known in the art, a 1's complement for a value written in binary or base-2 (i.e., has only zero's and one's) is the inverse of a existing one or zero. For example, a 1's compliment of  $110_2$  is  $001_2$ .

The ports-requested-field 82 is one-byte and has a variable value indicating a number of locally unique ports requested by a network device. By default the ports-requested-field 82 is 16 or 32, which is a reasonable number for most network devices. However, other default numbers could also be used. Unused-field 84 is three-bytes and has a value of zero. However, other layouts, values and field sizes could also be used for the PAP request message 66.

In one preferred embodiment of the present invention, a network device transmits a PAP request message 66 upon boot. The PAP 64 is associated with Dynamic Host Configuration Protocol ("DHCP") or BOOTstrap Protocol ("BOOTP"). DHCP is a protocol for passing configuration information such as IP 48 addresses to hosts on an IP 48 network. For more information on DHCP see RFC-1541 and RFC-2131, incorporated herein by reference. The format of DHCP messages is based on the format of BOOTP messages described in RFC-951 and RFC-1542, incorporated herein by reference. From a network device's point of view, DHCP is an extension of the BOOTP mechanism.

In another embodiment of the present invention, the network devices (14, 16, 18, 20, 22, 24) request locally unique ports after boot when a protocol layer in the layered protocol stack 42 makes an initial request for an external network (e.g., 30 or 32). The network devices (14, 16, 18, 20, 22, and 24) may also request more locally unique ports when the number of locally unique ports required falls below the number of locally unique ports allocated to the network devices.



The PAP request message 66 is sent from a network device (14, 16, 18, 20, 22, and 24) to the router 26 after attaching an IP 48 header or other message header. A PAP response message 68 is sent from the router 26 back to the network devices (14, 16, 18, 20, 22, 24) either confirming or denying the PAP request message 66.

FIG. 5 is a block diagram illustrating a PAP response message layout 86. A type-field 88 is one-byte and has a value for receiving responses (e.g., 32). A code-field 90 is one-byte and has a value of zero for failure and one for success. A checksum-field 92 is two-bytes and is a 16-bit 1's complement sum of the entire PAP response message 68. A lowest-port-field 94 is two-bytes and is a lowest locally unique port number allocated in a block of locally unique ports. A total-ports-field 96 is one-byte and is the total number of locally unique ports allocated to the network device. An unused-field 98 is one-byte and has a value of zero. However, other layouts, values and field sizes could also be used for the PAP response message 68.

Upon receiving a successful PAP response message 68, a network device saves the block of locally unique ports that it may use. The locally unique ports are saved in a data structure with a flag-field indicating whether the locally unique port is allocated or unused. Table 1 is pseudo-code for an exemplary data structures to store locally unique port information. However, other data structures or layouts could also be used.

```
struct unique_ports
{
    int port_number;
    flag status:1; /* one bit flag, 0 = unused, 1 = allocated */
} u_ports[MAX_GL];
int number_of_u_ports; /* number of locally unique ports allocated */
```

Table 1.

The one or more locally unique ports are allocated to protocols and applications in the layered protocol stack 42 on a network device to replace default or ephemeral ports. Upon receiving an unsuccessful PAP response message 68 a network device may send another PAP request message 66 for fewer ports. If the router 26 cannot allocate a large enough block of contiguous locally unique ports for the network device, it may send a PAP response 68 with a success code, but allocate fewer locally unique ports than requested.

FIG. 6 is a block diagram illustrating a PAP invalidate message layout 100. A PAP invalidate message 70 is used to invalidate or de-allocate a block of locally unique ports currently allocated to a network device. A type-field 102 is one-byte and has a value to de-allocate ports (e.g., 32). A code-field 104 is one-byte and has a value of two. A checksum-field 106 is two-bytes and is a 1's complement sum of the entire PAP invalidate message 70. A port-field 108 is one-byte and has a value of a locally unique port number used by the network device that is being invalidated or de-allocated. An unused-field 110 is three-bytes and has a value of zero. However, other layouts, values and field sizes could also be used for PAP invalidate message 70.

It is possible that two network devices may be allocated overlapping blocks of locally unique ports as a result of the router 26 crashing or rebooting. The router 26 should send a PAP invalidate message 70 to invalidate all locally unique ports in use upon reboot to help prevent this problem. A network device (14, 16, 18, 20, 22, and 24) also sends a PAP invalidate message 70 when it no longer needs a locally unique port.

FIG. 7 is a block diagram illustrating a combined network address layout 112 for combined network address 72. However, other layouts could also be used. The combined

network address layout 112 includes a common external network address 114 such as an IP 48 address (e.g., a common network address 28), and a locally-unique port 116 obtained by sending a PAP request message 66 and receiving a PAP response message 68 from a network device. The network devices (14, 16, 18, 20, 22, 24) use the combined network address 72 for

5 communications with the external second network 30 or the third network 32. The common external network address 114 identifies the first computer network 12 to an external second computer network (e.g., 30 or 32).

As is known in the art, to identify separate data streams, TCP 58 provides a source port field in a TCP 58 header and a source address field in an IP 48 header. For more information on

10 TCP headers see RFC-793. Since default or ephemeral port identifiers are typically assigned independently by a TCP 58 stack in a network, they are typically not unique. To provide for unique addresses within a TCP 58 stack, a local Internet address identifying a TCP stack 58 can be concatenated with a default or ephemeral port identifier, a remote Internet address and a remote port identifier to create an "association." The association is unique throughout all

15 networks connected together. Associations are known to those skilled in the networking arts.

In a preferred embodiment of the present invention, the source port in a header is given a locally unique port obtained with PAP 64 and given a common external network address. Together they uniquely identify applications and protocols on the network devices (14, 16, 18, 20, 22, 24) on the first computer network 12 to the second external computer network (e.g., 30 or

20 32) with a value conceptually similar to an association used by a TCP stack 58.

As is also known in the art, UDP 60 also has a source port field in a UDP header. For more information on UDP 60 headers see RFC-768. The UDP 60 source port is a non-optional

field. It indicates a port of the sending process and is assumed to be the port to which a reply should be addressed in the absence of any other information. If not used, a value of zero is inserted. A UDP 60 header also has a source address field. A locally unique port can also be used in a UDP 60 header.

5 In a preferred embodiment of the present invention, the PAP 64 is used to create combination network address 72 that is used in the TCP 58 or UDP 60 header fields. In another embodiment of the present invention, the combination network address 72 is stored in other message header fields understood by the router 26 (i.e., non-IP 48 TCP 58 or UDP 60 fields), the first computer network 12, the second computer network 30 and the third computer network 32.

10 In a preferred embodiment of the present invention, the router 26 allocates blocks of locally unique ports to network devices (14, 16, 18, 20, 22, and 24). However, other network devices could also be used to allocate locally unique ports (e.g., a port server). The router 26 maintains a port-to-internal network address table as locally unique ports are allocated. The router 26 also has an internal table indicating internal network addresses for all the network  
15 devices (14, 16, 18, 20, 22, 24) on the first computer network 12. In a preferred embodiment of the present invention, the internal network addresses for the first computer network 12 are private IP 48 addresses. For example, the computer 14 has an internal IP address of 10.0.0.1 (FIG. 1), the printer 16, 10.0.0.2, the computer 18, 10.0.0.3, the hand held computer, 20, 10.0.0.4, the telephone 22, 10.0.0.5, the facsimile, 24, 10.0.0.6, and the router 26, 10.0.0.7, in  
20 FIG. 1. The internal addresses are not published on the external computer network (e.g., the Internet or an intranet). However, other internal network addresses could also be used (e.g., Medium Access Control ("MAC") protocol addresses).

FIG. 8 is a block diagram illustrating a port-to-internal address table 118 layout maintained by the router 26. However, other layouts and more or fewer rows and columns could also be used. The port-to-internal address table 118 layout has three columns: an internal-network-address column 120, a lowest-port column 122, and a number-of-ports column 124. However, more or fewer columns or other table layouts could also be used. First row 126 indicates that a network device has been allocated ports 1026-1057 for use with internal network address, 10.0.0.1, (e.g., computer 14). A second network device has been allocated ports 1058-1073 for use with internal network address 10.0.0.3 (e.g., computer 18). An internal network address may have several entries in port-to-internal address table 118.

## **Distributed Network Address Translation**

FIG. 9 is a flow diagram illustrating a Method 130 for allowing distributed network address translation. At Step 132, a first network device on a first computer network requests one or more locally unique ports from a second network device on the first computer network with a first protocol. The locally unique ports are used to replace default or ephemeral ports in protocol layers in the layered protocol stack 42 on the first network device. In addition, the locally unique ports are used to create a combination network address 72 comprising a locally unique port and a common external address to communicate with a second external computer network without address translation. At Step 134, the first network device receives the one or more locally unique ports from the second network device. At Step 136, the first network device replaces one or more default or ephemeral ports used in the layered protocol stack 42 with one or more locally unique ports. At Step 138, the first network device constructs one or more combination network

addresses 72 using the one or more locally unique ports and a common external network address used to identify the first computer network on the second external computer network.

In a preferred embodiment of the present invention, the first network device is any of network devices (14, 16, 18, 20, 22, 24), the second network device is router 26, the first computer network is first computer network 12 (e.g., SOHO LAN) the first protocol is PAP 64, the second external computer network is any of the second computer network 30 (e.g., the Internet or an intranet) or the third computer network 32 (e.g., PSTN). The combination network address 72 includes a common IP 48 address (e.g., common network address 28) identifying network devices on the first computer network 12 to a second external computer network (e.g., 30 or 32). However, the present invention is not limited to the networks, network devices, network addresses or protocols described and others may also be used.

The locally unique ports are used for entities such as protocols and applications in layered protocol stack 42 on a network device and are locally unique on the first computer network 12. The locally unique ports will identify a network device on the first computer network 12. For example, TCP 58 typically has a default port or ephemeral port assigned to the TCP 58 stack (e.g., 1234). After allocation with Method 130, a network device uses a locally unique port to replace a default or ephemeral port in a protocol layer in the layered protocol stack 42. As is illustrated in FIG. 8, the network device 14 with an internal IP 48 address, 10.0.0.1, is assigned thirty-two locally unique ports in the range of 1026-1057. The network device 14 may assign locally unique port-1032 to TCP 58 to use as a default or ephemeral port. An original default port or ephemeral for TCP 58 was 1234. The combination network address 112 illustrated in FIG. 7 is then assigned to TCP 58 on the network device 14 for communications with an external

network (e.g., 30 or 32). Other locally unique ports are assigned to other protocols and applications in the layered protocol stack 42 on a network device to replace other default ports.

In one embodiment of the present invention, locally unique ports are assigned to protocol layers in the layered protocol stack 42 when a network device boots. In another embodiment of the present invention, locally unique ports are assigned to protocol layers in a layered protocol stack when a protocol layer makes a request for an external network (e.g., 30 or 32). In yet another embodiment of the present invention, locally unique ports are assigned dynamically or on-the-fly in an individual protocol layer as a protocol layer makes a request for an external network (e.g., 30 or 32).

The locally unique ports with common external network address 28 as the combination network address 112 uniquely identify an entity on a network device to an external network (e.g., 30 or 32) without translation. Network interface card device drivers 44 maintain the actual internal IP 48 address of a network device.

Locally unique-ports can also be used with the common external network address 28 (e.g., for Mobile IP). Locally unique ports help identify a mobile network device that roams away from a home network (e.g., first computer network 12) to a foreign network.

FIG. 10 is a flow diagram illustrating a Method 140 for distributed network address translation. At Step 142, a request is sent from a first network device on a first computer network to a second network device on the first computer network. The request is for a second external network and includes a combination network address 72 identifying the first network device on the first network. The combination network address 72 is constructed with Method 130 (FIG. 9) and includes a locally unique port and a common external address to identify the

first computer network to the second external network. At Step 144, the second network device routes the request from the first computer network to the second external network. At Step 146, the second network device on the first computer network receives a response from the external second computer network at the external network address identifying the first network from the combination network address. At Step 148, the second network device on the first computer network routes the response to the first network device on the first computer network using the locally unique port from the combination network address to identify the first network device.

In a preferred embodiment of the present invention, the first network device is any of network devices (14, 16, 18, 20, 22, 24), the second network device is router 26. The first computer network is first computer network 12, and the second computer network is second computer network 30 or third computer network 32. The combination network address includes a locally unique port obtained with PAP 64 and an external IP 48 address for an external network such as the Internet, an intranet, or another computer network. However, the present invention is not limited to the networks, network devices, network address or protocol described and others may also be used.

Method 140 (FIG.10) is illustrated with a specific example using TCP 58/IP 48 layers from the layered protocol stack 42. However, other protocol layers in the layered protocol stack 42 could also be used. At Step 142, the network device 14 sends a TCP 58 request to the server 39 (FIG. 1). For example, a TCP 58 request for server 39 at external IP 48 address, 192.200.20.3, on the second computer network 30. Table 2 illustrates an exemplary request data packet sent at Step 142.



IP 48 Header	TCP 58 Header
SRC IP: 198.10.20.30	SRC Port: 1032
DST IP: 192.200.20.3	DST Port: 80

Table 2.

The source IP 48 address is common external network address 28 (e.g., 198.10.20.30) and the source port is a locally unique port-1032 obtained via the PAP 64 with Method 130 and available to a TCP 58 service. In one embodiment of the present invention, the locally unique port-1032 replaces default port 1234 for TCP 58 when network device 14 was booted. In another embodiment of the present invention, default port 1234 is replaced with a locally a unique port, such as locally unique port-1032, whenever a protocol layer in layered protocol stack makes the request. The locally unique port along with the common external address comprise combination network address 112.

In one preferred embodiment of the present invention, the default TCP 58 port of 1234 has been replaced with a locally unique port-1032. The destination IP address is, 192.200.20.3, for the server 39 (FIG. 1) on the second external network 30 and the destination port is well known Internet port 80. When the request reaches a network interface card device driver 44 in the layered protocol stack 42, an outer IP 48 header is added to route the request to the router 26. For example, the outer IP 48 is a virtual tunnel header that is explained below. Network interface card device drivers maintain the local internal network address (e.g., 10.0.0.x) for a network device for internal communications. Table 3 illustrates an exemplary data packet with an outer IP 48 header added for router 26.

Outer IP 48 header	Inner IP 48 header	TCP 58 header
SRC IP: 10.0.0.1	SRC IP: 198.10.20.30	SRC Port: 1032
DST IP: 10.0.0.7	DST IP: 192.200.20.3	SRC Port: 80

Table 3.

A network interface card device driver 44 adds the outer IP 48 header including (e.g., a virtual tunnel header) a source IP 48 address for network device 14 of, 10.0.0.1, and a destination IP 48 address of, 10.0.0.7, for the router 26. At Step 144, the router 26 receives the request data packet, strips the outer IP 48 header, and sends the request data packet to the external network 30.

At Step 146, the router 26 receives a response packet from an external network (e.g., 30). An exemplary response data packet is illustrated in Table 4.

IP 48 Header	TCP 58 Header
SRC IP: 192.200.20.3	SRC Port: 80
DST IP: 198.10.20.30	DST Port: 1032

Table 4.

The router 26 receives the response packet from the external second network 30 at Step 146 with a destination IP 48 address for the common external network address, 198.10.20.30, and a destination port set to locally unique port-1032. The router 26 uses port-to-internal network address table (FIG. 8) to map destination port-1032 to an internal IP 48 address, 10.0.0.1, for the computer 14. The router 26 adds an outer IP 48 header (e.g., a virtual tunnel header) to route the response data packet sent back to the network device 14. Table 5 illustrates an exemplary response packet with an outer IP 48 header added by the router 26.

Outer IP 48 header	Inner IP 48 header	TCP 58 header
SRC IP: 10.0.0.7	SRC IP: 192.200.20.3	SRC Port: 80
DST IP: 10.0.0.1	DST IP: 198.10.20.30	DST Port: 1032

Table 5.

The outer IP 48 header has a source internal IP 48 address of, 10.0.0.7, for the router 26 and a destination internal IP 48 address of, 10.0.0.1, for the network device 14 on computer network 12. At Step 148, the router 26 routes the response data packet to the network device 14 with the outer IP 48 header. A network interface card device driver 44 in the layered protocol

stack 42 strips the outer IP 48 header and forwards the response data packet to the network layer 46. This step can also be done in the device driver.

The network device 14 sends a request to an external network and receives a response from the external network using DNAT and locally unique port-1032 allocated with the PAP 64.

5 The router 26 does not translate any source/destination IP 48 addresses or source/destination ports. Thus, DNAT is accomplished without NAT at the router 26.

10 A preferred embodiment of the present invention is described with respect to a single common external network address identifying multiple network devices on first computer network 12 and used in combination network address 112 with a locally unique port. However, the present invention is not limited to a single common external network address and can also be practiced with a multiple common external network addresses.

15 Distributed NAT using Method 130 (FIG. 9) and Method 132 (FIG. 10) removes the computation burden of NAT at the router 26 and allows multiple network devices to use a single or a small number of external network addresses known to an external network such as the Internet or an intranet. Instead of providing NAT, the router 26 routes data packets from a network device (14, 16, 18, 20, 22, 24) on the first computer network 12 to a second external computer network such as the second computer network 30 or the third computer network 32 using the combination network address. In addition, the router 26 is no longer required to support multiple application protocols from the layered protocol stack 42.

20 The router 26 also routes data packets from the second external computer network back to a network device on the first computer network using the locally unique port in the combination network address. The router 26 is no longer required to replace an internal network

address with an external network address for outbound traffic, and replace an external network address with an internal network address for inbound traffic. Thus, DNAT of the present invention removes the computational burden of NAT from the router 26 and does not violate the Internet principal of providing end-to-end transmission of data packets between network devices without alternations.

### DNAT with Port Translation

In another preferred embodiment of the present invention, DNAT is accomplished without modifying protocols or applications in the layered protocol stack 42 above the network interface device driver layer 44. However, in such an embodiment, a network interface card device driver 44 in the network devices (14, 16, 18, 20, 22, 24) is used to translate default or default ports on-the-fly to/from locally unique ports reserved by a network device with the PAP 64. In addition, the network interface card device driver 44 supports multiple protocols from the layered protocol stack 42 for DNAT with port translation.

As an example, suppose the computer 14 (FIG. 1) with an internal IP 48 address, 10.0.0.1, makes a TCP 58/IP 48 request from a server on the second computer network 32 (e.g., the Internet) at external IP 48 address, 192.200.20.3, (i.e., web server 39, FIG. 1). The initial TCP 58 packet reaching network interface card device driver 44 of layered protocol stack 42 is illustrated in Table 6.

IP 48 Header	TCP 58 Header
SRC IP 198.10.20.30	SRC Port: 1234
DST IP 192.200.20.3	DST Port: 80

Table 6.

The local source port for TCP 58 is 1234, the destination port is well known port 80 for the Internet, the source IP 48 address is the common external network address 28 and the destination address is external IP 48 address for server 39 (FIG. 1).

In the preferred embodiment discussed above using Methods 130 and 140 of FIGS. 9 and 10, application and/or protocol local default ports are modified by a network device to use a locally unique port obtained via the PAP 64 in protocol layers above the device drivers.

However, for DNAT with port translation, ports are not translated in the layered protocol stack 42. Network interface card device drivers instead provide port and address translation. In such an embodiment, a network interface card device driver 44 will determine that a connection is being initiated. An entry in a Source Port Translation Table ("SPTT") in a network interface card device driver 44 is created.

FIG. 11 illustrates a SPTT layout 150. However, other layouts, field sizes and values could also be used. A default-port field 152 is two-bytes and is a default or ephemeral port number used by a TCP 58 service and other applications of a network device. A translated-port 154 field is two-bytes and is a locally unique port number used for external communications allocated by PAP 64. A protocol-field 156 is one-byte and has a value of zero for TCP 58 and a value of one for UDP 60. A timestamp-field 158 is four-bytes and includes a value of a current system time in milliseconds updated every time this entry is used.

The TCP 58 source port, 1234, is translated into a locally unique port allocated by the PAP 64 by a network interface card device driver. The TCP 58 source port, 1234, is not translated in the TCP 58 layer or any other protocol layer above the network interface card

device driver 44 in the layered protocol stack 42. An entry is added to SPTT 150. Table 7 illustrates an exemplary SPTT 150 table entry.

Default Port	Locally Unique Port	Protocol	Timestamp
1234	1032	1 (TCP)	10023

Table 7.

After translation by the network interface card driver, an outer IP 48 header is added to the data packet. The outer IP header is used for routing (e.g., through a virtual tunnel). The outer IP header has the internal address of the network device as a source IP 48 address (e.g., 10.0.0.1) and the internal network address of router 26 (e.g., 10.0.0.7) as a destination address. Table 8 illustrates the data packet with the outer IP 48 header.

Outer IP 48 Header	Inner IP 48 Header	TCP 58 Header
SRC IP 10.0.0.1	SRC IP 198.10.20.30	SRC port 1032
DST IP 10.0.0.7	DST IP 192.200.20.3	DST port 80

Table 8.

Upon receiving the data packet illustrated in Table 4, the router 26 examines the source port (e.g., 1032) and the outer IP 48 source address (e.g., 10.0.0.1) to ensure a network device is using a valid locally unique port assigned to the network device. Router 26 maintains an IP Address Translation Table ("IPATT").

FIG. 12 illustrates an exemplary IPATT layout 160. However, other layouts, field sizes and values could also be used. A destination port-field 162 is two-bytes and holds a locally unique port obtained with PAP 64. An internal destination IP address-field 164 is four-bytes and is the internal IP 48 address (e.g., 10.0.0.1) of a network device using the locally unique port in destination port-field 162. A protocol-field 166 is one-byte and has a value of zero for TCP 58 or a value of one for UDP 60. A timestamp-field 168 is four-bytes and includes a value of a

current system time in milliseconds updated every time this entry is used. Table 9 illustrates an exemplary IPATT 160 table entry.

Destination Port (locally unique port)	Internal Destination IP 48 Address	Protocol	Timestamp
1032	10.0.0.1	6 (TCP)	10048

Table 9.

- 5 Table 9 illustrates a locally unique port-1032 is associated with internal IP 48 address 10.0.0.1 (e.g., computer 14) for the TCP 58 protocol. The router 26 strips off the outer IP 48 header illustrated in Table 4 and sends the data packet comprising the inner IP 48 header and TCP 58 header to the external network 30.

- 10 A response data packet arrives from an external network on common external network address 28 (e.g., 198.10.20.30). An arriving packet contains the headers illustrated in Table 10.

IP 48 Header	TCP Header
SRC IP 192.200.20.3	SRC Port: 80
DST IP 198.10.20.30	DST Port: 1032

Table 10.

- The router 26 looks up the destination port (i.e., locally unique port-1032) in IPATT 158 (Table 9) and finds local network address, 10.0.0.1, (e.g., for computer 14). The router 26 then creates an outer IP 48 header such as the exemplary IP 48 header illustrated in Table 11. The outer IP 15 48 header has a source IP 48 address for the router 26 and a destination IP 48 address for network device 14.

Outer IP 48 Header	Inner IP 48 Header	TCP 58 Header
SRC IP 10.0.0.7	SRC IP 192.200.20.3	SRC port 80
DST IP 10.0.0.1	DST IP 198.10.20.30	DST port 1032

Table 11.

- 20 The router 26 then transmits the data packet illustrated in Table 11 to the appropriate network device (e.g., computer 14 at internal address 10.0.0.1). Upon receiving the data packet,

a network interface card driver looks up the destination port (e.g., 1032) in the SPTT 148 (e.g., Table 7) finding a mapping to TCP 58, port 1234. The locally unique port-1032 is re-translated back to TCP 58 default port 1234 in the device driver. No translation is done above the device driver. The outer IP 48 header is then stripped. The data packet is forwarded to IP 48 in the network layer 46. Table 12 illustrates the forwarded data packet.

Inner IP 48 header	TCP 58 header
SRC IP 192.200.20.3	SRC Port 80
DST IP 198.10.20.30	DST Port 1234

Table 12.

The end of the connection is detected by both the router 26 and the network device 14. Upon end of connection, the entries in the SPTT 148 and IPATT 160 tables are removed from the router 26 and network interface card driver.

FIG. 13 illustrates a Method 170 for outbound distributed NAT using port translation. At Step 172, a network interface card device driver 44 receives a data packet from the network layer 46 (e.g., Table 6). At Step 174, the network interface card device driver 44 conducts a test to determine if a destination network address (e.g., 192.200.20.3) is for an external network (e.g., 30 or 32). If so, at Step 176, the network interface card device driver 44 adds an outer IP 48 header (e.g., a virtual tunnel header) to the data packet with the source address set to the network device's internal IP 48 address (e.g., 10.0.0.1) and the destination address set to the router's 26 internal address (e.g., 10.0.0.7) as (e.g., Table 8). At Step 178, a local source port for the application or protocol from the header (e.g., TCP 58 port 1234) is translated into a locally unique port (e.g., 1032) obtained via PAP 64 with SPTT 150 (e.g., Table 7). At Step 180, the data packet with the outer IP 48 header is transmitted to network interface card hardware, which forwards to data packet to the router 26.



If the test at Step 174 determines that the destination network address is for internal network 12, then at Step 182, the default or ephemeral source port is not translated to a locally unique port for internal communications. Using Method 170, distributed NAT is done by a network interface card device driver, and no port translation occurs above device driver.

- 5 However, other software or hardware modules or drivers besides network interface card device driver 44 could also translate ports with Method 170.

FIG. 14 is a flow diagram illustrating a Method 184 for inbound distributed NAT using port translation. At Step 186, a data packet is received on a network interface card driver 44 (e.g., Table 11) from the router 26. The router 26 received the data packet from external network 30 or 32 and added an outer IP 48 header. At Step 188, a test is conducted to determine if the source IP 48 address from the inner IP 48 header is an external IP 48 address. If so, at Step 190 the destination port from the inner IP 48 header is translated from a locally unique port to a default port (e.g., 1032 → 1234) using the SPTT 158 (Table 7). At Step 192, the outer IP 48 header is stripped off. At Step 192, the data packet (e.g., Table 12) is forwarded to the network layer 46.

If the test at Step 188 determines that the source IP 48 address is for the internal network 12, then at Step 196 the source IP 48 address from the outer IP 48 header is copied to the inner source IP 48 address. At Step 192, the outer IP 48 header is stripped off. At Step 194, the data packet is forwarded to network layer 46. The default or local source port is not translated to a locally unique port for internal communications.

Using Method 184, distributed NAT is done by a network interface card device driver, and no port translation occurs above the device driver. However, other software or hardware

modules or drivers besides a network interface card device driver, or in layers above the network interface card device driver 44 could also translate ports with Method 184.

DNAT (FIG. 9 and FIG 10) does port translation in individual protocol layers in the layered protocol stack 42. The port translation is done at boot time for a network device, or  
5 dynamically in a protocol layer when a protocol layer makes a request to an external network (e.g., 30 or 32).

In contrast, DNAT with port translation (FIG. 13 and FIG. 14) does port translation in the network interface card device driver 44 on a network device. No ports are translated in protocol layers above the device driver. In addition, the network interface card device driver 44 supports  
10 multiple protocols from the layered protocol stack 42 above the network interface card device driver 44 for DNAT with port translation. For outbound data, a default port assigned to an application or protocol is translated to a locally unique port "on-the-fly" in the device driver. For inbound data, the network device translates a locally unique port back to a default port on-the-fly in the device driver. DNAT with on-the-fly port translation in the network interface card device  
15 driver 44 (FIGS. 13 and 14) places more computational overhead on a network device than DNAT with port translation in individual protocol layers (FIG. 10).

However, DNAT with on-the-fly port translation in the network interface card device driver 44 (FIGS. 13 and 14) is still preferred over non-distributed NAT in the router 26 with  
20 Methods known in the art since computational costs for translation are distributed among a number of network devices and not concentrated in the router 26. The router 26 does not translate any addresses for the described embodiments of the present invention. The method and

protocol for distributed NAT described above can also be used with protocols that provide security for a network using IP 48.

### Internet Protocol Security

There are a number of security measures that can be used with IP 48. One or more security measures can be indicated in an IP 48 header. IPSEC processing is confined completely within the IP 48 layer. All DNAT processing, when used with IPSEC must run above the IP 48 layer. Otherwise, IPSEC parameters are violated.

FIG. 15 is a block diagram illustrating an IP 48 packet header 200. A version-field 202 includes an IP 48 protocol version (e.g., IPv4 or IPv6). An Internet Header Length ("IHL")-field 204 includes a length for the header. A Type-of-Service ("ToS")-field 206 includes a requested type of service. A total length-field 208 includes a length of everything in an IP 48 data packet including the IP 48 header 200. An identification-field 210 is used with packet fragmentation. A fragment offset field 212 is also used with packet fragmentation. A Time-To-Live ("TTL")-field 214 is now a hop count used to limit a lifetime for an IP 48 packet included with the header. A protocol-field 216 includes a protocol used with the IP 48 packet 200 (e.g., TCP 58, UDP 60, ESP, AH, etc.). A header checksum-field 218 is used to verify the contents of the IP 48 packet header 200. A source address-field 220 includes a source IP 48 address for a sending endpoint. A destination address-field 222 includes an IP 48 address for a receiving endpoint. An options-field 224 is used for security, source routing, error reporting, debugging, time stamping, and other information. IP 48 data (e.g., TCP 58, UDP 60, etc.) appears below the options-field 224.

IPSEC provides security for IP 48 packets. For more information in IPSEC see "Security Architecture for the Internet Protocol", by S. Kent and R. Atkinson, RFC-2401, November,

1998, incorporated herein by reference. Three security requirements are typically addressed by IPSEC. IPSEC provides message authentication, integrity and confidentiality for IP 48 packets moving between a source and a destination endpoint. Starting from a state in which no connection exists between two endpoints, a Security Association ("SA") can be established  
5 based upon IP 48 such that each endpoint trusts the security of the connection, and an identity of each endpoint is authenticated to the other.

IPSEC typically defines two security services, each having an associated header that is added to an IP 48 packet that it protects. The two security services are an Authentication Header ("AH") and an Encapsulating Security Payload ("ESP") header. However, more or fewer  
10 security services can also be used with IPSEC.

The AH provides authentication and integrity protection for IP 48 packets. For more information on the AH see, "IP Authentication Header," by S. Kent and R. Atkinson, RFC-2402, November, 1998, incorporated herein by reference.

The ESP provides encryption protection as well as optional authentication and integrity  
15 protection. For more information on the ESP see, "IP Encapsulating Security Payload (ESP)," by S. Kent and R. Atkinson, RFC-2406, November, 1998, incorporated herein by reference.

The IPSEC protocol headers are identified in the protocol-field 216 of an IP packet header 200 (FIG. 15). An IPSEC protocol header specifies a protocol type (i.e., AH or ESP) and contains a numerical value called the Security Parameter Index ("SPI"). The SPI is a unique  
20 identifier associated with a SA by a receiving endpoint. The identifying information is used by a receiving endpoint to help it correctly associate an IP 48 packet with a SA. Correct association of an IP 48 packet with a SA is required in order to apply proper IPSEC processing.

The IPSEC services can be applied in one of two modes, a "transport mode" or a "tunnel mode." In the transport mode, a packet is routed directly to its final destination according to a destination address (e.g., IP 48 destination address 222 (FIG. 15)). A final destination is where the IPSEC processing is done, as well as where the IP 48 packet is "consumed," (i.e., processed).

5 The destination IP 48 address is "visible" (i.e., not encrypted) as the IP 48 packet traverses the network.

As is known in the art, a virtual tunnel can be created by encapsulating a data packet inside another data packet. For example, an outer header is added before an inner header of a data packet (e.g., Tables 3, 5, 8 and 11). Between the inner header and outer headers are any

10 other headers for a data path, or security, such as security headers specific to a tunnel configuration. The outer header typically identifies the "endpoints" of the tunnel. The inner header typically identifies an original sender and recipient of the data. For more information, see "IP-in-IP tunneling," by W. Simpson, RFC-1853, October 1995, incorporated herein by reference.

15 In the tunnel mode, an outermost tunnel IP 48 header encapsulates a protected IP packet. A first destination address is an endpoint of a tunnel according to a tunnel destination address. A final destination address is not necessarily the same as an endpoint address of the tunnel. A destination IP 48 address 222 (FIG. 15) in the IP 48 header of the encapsulated (i.e., encrypted) part may or may not be "visible."

20 IPSEC protocols establish and use a Security Association ("SA") to identify a secure virtual connection between two endpoints. A SA is a unidirectional connection between two endpoints that represents a single IPSEC protocol-mode combination. Two termination

endpoints (i.e., network devices for the transport mode, or intermediate devices for the tunnel mode) of a single SA define a secure virtual connection that is protected by IPSEC services. One of the endpoints sends IP 48 packets, and the other endpoint receives them. Since a SA is unidirectional, a minimum of two SAs are required for secure, bi-directional communications. It is also possible to configure multiple layers of IPSEC protocols between two endpoints by combining multiple SAs.

FIG. 16 is a block diagram illustrating an IPSEC Authentication Header 226. A next header-field 228 is an 8-bit field that identifies the type of the next payload after the AH. A payload length-field 230 specifies the value of an AH in 32-bit words (i.e., 4-bytes). A reserved-field 232 is a 16-bit field reserved for future use. A Security Parameters Index ("SPI")-field 234 is an arbitrary 32-bit value that, in combination with a destination IP 48 address and a security protocol (e.g. AH or ESP), uniquely identify a SA for the data packet. A set of SPI values are in the range of 1 through 255 are reserved by the Internet Corporation for Assigned Names and Numbers ("ICANN") for future use. More information on ICANN can be found at the URL "www.icann.org." A SPI greater than 255 is selected by a destination endpoint upon establishment of a SA. Allocation of SPI using the PAP 64 is explained below. A sequence number-field 236 is an unsigned 32-bit field including a monotonically increasing counter value as a sequence number. An authentication data-field 238 is a variable length field that contains an Integrity Check Value ("ICV") for a packet.

In the transport mode, a sending endpoint inserts an AH header after an IP 48 header and before an upper protocol layer (e.g., TCP 58, UDP 60, etc.). In the tunnel mode, outer and inner IP header/extensions can be used in a variety of ways. Placement of the AH header in the tunnel

mode is dependent on a variety of factors including the type of tunneling used. Thus, a location for an AH header may vary.

For outbound packets, AH is applied after an IPSEC application determines that a packet associated with a SA wants AH processing. A sending endpoint's AH sequence number-field 236 (FIG. 16) is initialized to zero when a SA is established. The sending endpoint increments the sequence number-field 236 for a SA. Thus, a first AH packet using a given SA will have a sequence number of 1. An AH ICV used in the authentication data-field 238 (FIG. 16) is computed over IP header fields 200 (FIG. 15) that are either immutable in transit, or are predictable in value upon arrival at an endpoint for the AH SA. The AH header 226 (FIG. 16) and explicit padding bytes, if any, are computed after the IP 48 header 200 fields (FIG. 15). Upper level protocol data (e.g., TCP 58, UDP 60), which is assumed to be immutable in transit is computed last. If required, IP 48 fragmentation occurs after AH processing using an IPSEC implementation.

For inbound packets, packet reassembly is performed prior to AH processing. Upon receipt of a packet containing an AH, a receiving endpoint determines an appropriate SA, based on a destination IP 48 address 222 (FIG. 15), a AH protocol header 226 (FIG. 16), and an AH SPI 234 (FIG. 16). A sequence number is verified next. The sequence number helps prevent replay attacks. An ICV value is computed over appropriate fields of the packet, using a specified authentication algorithm, and verifies that it is the same algorithm as the ICV included in the authentication data-field 238 of the AH header 226 (FIG. 16).

FIG. 17 is a block diagram illustrating an ESP packet format 240. A SPI-field 242 is an arbitrary 32-bit value that, in combination with a destination IP 48 address and a security

protocol (e.g. AH or ESP), uniquely identify a SA for the data packet. A sequence number-field 244 is a 32-bit field that includes a monotonically increasing counter value as a sequence number. A payload data-field 246 is a variable length field including data described by the next header field 248. A padding-field 250 is used with the payload data-field 246 for encryption. A pad length-field 252 indicates a number of pad bytes immediately preceding it. A next header-field 248 is an 8-bit field that includes a type of data contained in the payload data-field 246. An authentication data-field 254 is a variable length field including an Integrity Check Value ("ICV") computed over the whole ESP header 240 minus the authentication data-field 254.

In the transport mode, a sending endpoint encapsulates upper layer protocol information in an ESP header and trailer and retains an original IP 48 header. In the tunnel mode, the outer and inner IP 48 headers/extensions can be inter-related in a variety of ways depending on the encryption being used. Thus, a location for the ESP may vary.

For outbound packets, ESP is applied after an IPSEC application determines that a packet associated with a SA wants ESP processing. The sending endpoint encapsulates into the ESP payload data-field 246 (FIG. 17) and original upper layer protocol information for the transport mode using a selected encryption technique. An entire IP 48 data packet is encapsulated for the tunnel mode. Any necessary padding is added to the padding-field 250. The payload data-field 246, the next header-field 248, the padding-field 250, and the padding length-field 252 are encrypted with an encryption technique. The exact steps used for constructing an outer IP 48 header depend on the mode (e.g., transport or tunnel) and the encryption technique being used.

A sending endpoint's sequence number-field 244 is initialized to zero when a SA is established. The sending endpoint increments the sequence number field 244 for a SA. Thus, a



first ESP packet using a given SA will have a sequence number of 1. If authentication is selected for the SA, the sending endpoint computes an ICV over the whole ESP header 240 minus the authentication data-field 254. If necessary, fragmentation is performed after ESP processing with an IPSEC implementation.

5 For inbound packets, packet reassembly is performed prior to ESP processing, if necessary. Upon receipt of an IP 48 packet including an ESP header 240, a receiving endpoint determines the appropriate SA based on a destination IP address 222 (FIG. 15), ESP protocol header 240 (FIG. 17), and a SPI 242 (FIG. 17). The SA indicates whether the sequence number-field 244 will be checked, whether the authentication data-field 254 should be present, and what encryption techniques should be used for decryption and ICV computations, if necessary.

10 During decryption, the ESP payload data-field 246, next header-field 248, the padding-field 250, and the padding length-field 252 are decrypted using a key, decryption technique, and cryptographic synchronization data if any, indicated by the SA. Any padding from the padding-field 250 is processed if necessary. An original IP 48 packet is reconstructed including an

15 original IP 48 header 200 (FIG. 15) plus original upper layer protocol information for the transport mode in the ESP payload data-field 246 (FIG. 17). A tunnel IP 48 header and an entire IP 48 packet is reconstructed in the ESP payload data-field 246 for the tunnel mode. The exact steps for reconstructing the original IP 48 packet depend on the mode (i.e., transport or tunnel).

FIG. 18 is a block diagram illustrating end-to-end security 256 between two endpoints

20 across an IP 48 network 30 (e.g., the Internet or an intranet) using AH, ESP and combinations thereof, in the transport and tunnel modes. A first end point 258, has a secure connection 260 to a second endpoint 262. A first exemplary data packet 264 includes a first IP 48 address ("IP1")

in a first IP 48 header, an AH header and upper level protocol data. A second exemplary data packet 266 includes a first IP 48 address, an ESP header and upper level protocol data. A third exemplary data packet 268 includes a first IP 48 address, an AH header, an ESP header, and upper level protocol data. The exemplary data packets 264, 266 and 268 are used in the transport mode. One type of data packet layouts is typically selected (264, 266, or 268) for the transport mode depending on the type of security desired.

In the tunnel mode, a fourth exemplary data packet 270 includes a tunnel IP 48 header with a tunnel IP address ("TIP"), an AH header, an original IP 48 header with a first IP 48 address ("IP1") and upper level protocol data. A fifth exemplary data packet 272 includes a tunnel IP 48 header with a tunnel IP 48 address, an AH header, an original IP 48 header with a first IP 48 address and upper level protocol data. One type of exemplary data packet 270 or 272 is typically selected for the tunnel mode depending on the security desired. A combination of AH and ESP in the tunnel mode is not typically used and is not illustrated in FIG. 18. However, a combination of AH and ESP may be also be used in the tunnel mode with the present invention.

A set of protocols has been developed to allow two endpoints to establish one or more SAs between them. The process of establishing an IPSEC SA involves both negotiation and authentication. The negotiation results in an agreement between the two endpoints as to which security protocol and mode to use, as well as specific encryption techniques, associated parameter values, and SPI assignment for each SA that was established. The authentication ensures that each endpoint can trust the identity of the other endpoint during negotiation, and hence after the SA is established.

A number of standards have been proposed for protocols that establish SAs including an Internet Security Association and Key Exchange Protocol ("ISAKMP"), an Oakley Protocol ("Oakley"), and the Internet Key Exchange ("IKE") protocol, which incorporates ISAKMP and Oakley. For more information on ISAKMP see, "Internet Security Association and Key Management Protocol ("ISAKMP")," by D. Maughan, M.Schertler, M. Schneider and J. Turner, RFC-2408, November, 1998, incorporated by reference. For more information on Oakley see, "The OAKLEY Key Determination Protocol," by H.K. Orman, RFC-2412, November, 1998, incorporated herein by reference. For more information on IKE see, "The Internet Key Exchange (IKE)," by D. Harkins and D. Carrel, RFC-2409, November, 1998, incorporated herein by reference.

Using ISAMKP and IKE, SA negotiation is carried out as a sequence of signaling exchanges between two endpoints. A first endpoint proposes a security protocol and encryption algorithm, and a second endpoint accepts or counter-proposes. Once the signaling is complete both endpoints have agreed to negotiated details, relevant security parameter information is exchanged and the endpoints are ready to send or receive on a single unidirectional SA. Part of the signaling includes exchange of authentication information, using a CA. This is described below.

Authentication is based on a trusted third-party called a Certificate Authority ("CA"). Each endpoint that participates in IPSEC generates a public/private encryption key pair, and has its public key "notarized" by the CA. The CA binds an endpoint's IP 48 address to its public key, generates a certificate and returns it to an owner of the key. Thus, IP 48 addresses are one "security name space" used for binding public keys to their owners.

During SA negotiation, one endpoint supplies another endpoint with its certificate along with a signature that is encrypted with its private key. The certificate and signature are verified with a public key. A recipient (one at each endpoint) uses a sender's public key from its certificate to validate the signature and the sender's right to use its IP 48 address. Since only the sender has access to the private key, the recipient, once it has verified the signature, is certain of the initiator's "identity." In one exemplary preferred embodiment of the present invention, the identity is determined by the IP 48 address of the initiator, as IP 48 addresses form the security name space used to bind public keys to their owners. However, other security name spaces could also be used using other than an IP 48 address for an initiator's identity. Certificates are issued with a "Time-to-Live" value, after which they expire and become invalid. The result of negotiation and authentication is a secure connection 260 (FIG. 18) for one unidirectional SA. A second SA for bi-directional communications may be registered in a similar manner.

As was discussed above, NAT routers known in the art need to modify IP 48 packets. However, once an IP 48 packet is protected by IPSEC, it cannot be modified anywhere along its path to the IPSEC destination. NAT routers known in the art typically violate IPSEC by modifying packets. In addition, even if a NAT router did not need to modify the packets it forwards, it must be able to read the TCP 58 or UDP 60 port numbers. If ESP is used by a local endpoint, the port numbers will be encrypted, so the NAT router will not be able to complete its required mapping.

Local network devices on a LAN that use NAT possess only local, non-unique IP 48 addresses. These do not comprise a security name space that is suitable for binding a public key to a unique identity (i.e., a unique global IP 48 address). Without this binding, it is typically not

possible to provide the authentication necessary for establishment of SAs. Without authentication, neither endpoint can be certain of the identity of their counter part, and thus cannot establish a secure and trusted connection via a SA. However, DNAT described above, can be used with IPSEC to overcome some of the problems with NAT devices known in the art.

## 5 **Distributed Network Address Translation and IP Security**

A network device using DNAT as described above may also desire to establish a secure virtual connection to an external network device using IPSEC (e.g., SPIs). Such a network device would request and use locally unique ports and use DNAT as was described above. In addition, the network device may request locally unique security values to use DNAT with IPSEC.

FIG. 19 is a flow diagram illustrating a Method 274 for distributed NAT with security.

At Step 276, a first network device on a first computer network requests with a first protocol, one or more locally unique security values (e.g., SPIs) from a second network device on the first computer network and for distributed NAT. The one or more locally unique security values are used to identify security associations for data reception on the first network device during secure communications with a third network device on a second external network. At Step 278, the one or more locally unique security values are received on the first network device from the second network device with the first protocol. The one or more locally unique security values are stored on the first network device at Step 280. The one or more locally unique security values can be used to identify a unique security association for secure communications and used for distributed NAT. A unique security association identified by the first computer on the first network is used for reception of packets on the first computer.

In one exemplary preferred embodiment of the present invention, the first network device is a network device (14, 16, 18, 20, 22, and 24), the second network device is the router 26, the first protocol is the PAP 64, the one or more locally unique security values are SPIs used with IPSEC, including AH or ESP. In one exemplary preferred embodiment of the present invention, the locally unique security values are obtained with the PAP 64 using a PAP 64 security request message 67, a PAP 64 security response message 69, and a PAP 64 security invalidate message 71.

FIGS. 4A, 5A, and 6A illustrate exemplary PAP 64 security request message 67 layout 73, a PAP 64 security response message 69 layout 87, and a PAP 64 security invalidate message 71 layout 99. The PAP 64 security messages are used to allocate and de-allocate locally unique security values (e.g., SPIs) and are similar to the PAP 64 messages used to allocate locally unique security values.

FIG. 4A is a block diagram illustrating a PAP security request message 67 layout 73. A type-field 75 is one-byte and has a value (e.g., 33) for requesting locally unique security values. A code-field 77 is one-byte and has a value of zero for locally unique security values. A checksum-field 79 is two-bytes, and has a value of a 1's complement sum of the entire PAP security request message layout 73. The security values-requested-field 81 is two-bytes and has a variable value indicating a number of locally unique security values requested by a network device. Unused-field 83 is two-bytes and has a value of zero. However, other layouts, values and field sizes could also be used for the PAP security request 67 message layout 73.

FIG. 5A is a block diagram illustrating a PAP security response message 69 layout 85. A type-field 87 is one-byte and has a value for receiving security responses (e.g., 33). A code-

field 89 is one-byte and has a value of zero for failure and one for success. A checksum-field 91 is two-bytes and is a 16-bit 1's complement sum of the entire PAP security response message 85. A total-security-value-field 93 is two-bytes and is the total number of locally unique ports allocated to the network device. An unused-field 95 is two-bytes and has a value of zero. A lowest-unique-security-value-field 97 is four-bytes and includes a lowest locally unique security value allocated in a block of locally unique security values. However, other layouts, values and field sizes could also be used for the PAP security response message 85.

FIG. 6A is a block diagram illustrating a PAP security invalidate message 71 layout 99. A type-field 101 is one-byte and has a value to de-allocate security values (e.g., 33). A code-field 103 is one-byte and has a value of two. A checksum-field 105 is two-bytes and is a 1's complement sum of the entire PAP security invalidate message 99. A security-value-field 107 is four-bytes and has a value of a locally unique security value used by the network device that is being invalidated or de-allocated. However, other layouts, values and field sizes could also be used for PAP security invalidate message 99.

Returning to FIG. 19, the first network device, such as a computer 14, uses a PAP 64 security request message 67 to request the locally unique SPIs, and receives the SPIs in a PAP 64 security response message 69. The locally unique SPIs are requested, received and stored in a manner similar to the locally unique DNAT ports described above. However, the present invention is not limited to this exemplary preferred embodiment, and other network devices, protocols and security values could also be used. In one exemplary preferred embodiment of the present invention, the second network device allocates the one or more locally unique security values used on the first network device.

FIG. 20 is a flow diagram illustrating a Method 282 for distributed NAT with security.

At Step 284, a request message in a first protocol is received on a second network device requesting one or more locally unique security values for a first network device. At Step 286, one or more locally unique security values are allocated on the second network device. At Step 288, a network address for the first network device is stored with the one or more locally unique security values in a table associated with the second network device. The table is used to maintain a mapping between a network device and a locally unique security value for distributed NAT with security. At Step 290, the one or more locally unique security values are sent in a response message with the first protocol to the first network device.

In one exemplary preferred embodiment of the present invention, the first network device is a network device (14, 16, 18, 20, 22, 24) on the first computer network 12, the second network device is the router 26, the first protocol is PAP 64, the one or more locally unique security values are SPIs used with IPSEC including AH or ESP. The first network device, such as customer computer 14, uses a PAP 64 security request message 67 to request the locally unique SPIs. At Step 284 (FIG. 20), the router 26 receives the PAP 64 security request message 67. The router 26 maintains a table similar to the port-to-internal-network address table 118 illustrated in FIG. 8 except that a SPI value is used in place of a port number. At Step 286, the router 26 allocates one or more locally unique SPIs. At Step 288, a local IP 48 address for the first network device (e.g., 10.0.0.1) is stored with the one or more locally unique SPI values in a table associated with the second network device (e.g., see FIG. 21 below). The table is used to maintain a mapping between a network device and a locally unique SPI for distributed NAT with security. At Step 290, the one or more locally unique SPIs are sent by the router 26 in a PAP 64



security response message 69 to the first network device 14. However, the present invention is not limited to this exemplary preferred embodiment, and other network devices, protocols, messages, tables and security values could also be used with Method 282.

FIG. 21 is a block diagram illustrating a SPI-to-internal network address table layout 292 used at Step 288 of Method 284 (FIG. 20). FIG. 21 is similar to FIG. 8 except that the locally unique SPI values are 32-bits and the locally unique port values are 16-bits. In FIG. 21, an internal network address column 294 includes internal network addresses for network devices (14, 16, 18, 20, 22, 24) on the first computer network 12. The lowest SPI column 296 includes a lowest SPI value allocated. The number of SPIs column 298 includes a total number of locally unique SPIs allocated to a network device. For example, at row 300, a first network device 14 (FIG. 1) with a local IP 48 address of 10.0.0.1 on the first computer network 12, has been allocated 32 SPIs beginning with a SPI of value "280." At row 302, another network device 18 with a local IP 48 address of 10.0.0.3 on the first computer network 12, has been allocated 16 SPIs beginning with a SPI value of "312." However, the present invention is not limited to this SPI-to-internal network address table layout, and other SPI-to-internal network address table layouts can also be used. A first network device (e.g., 14, FIG. 1) will use locally unique security values (i.e., SPIs) with a second secure protocol (e.g., IPSEC) to establish a virtual secure connection (i.e., a SA) to a third external network device (e.g., 39 FIG. 1).

#### **Establishing IPSEC security associations using DNAT**

As was discussed above, the process of establishing an IPSEC SA involves both negotiation and authentication. Authentication is based on a trusted third-party called a Certificate Authority ("CA"). Each endpoint that participates in an IPSEC SA generates a

public/private encryption key pair, and has its public key “notarized” by the CA. The CA binds an endpoint’s IP 48 address to its public key, generates a certificate and returns it to an owner of the key. Thus, IP 48 addresses are used to provide a name space for binding public keys to their owners.

5 In one exemplary preferred embodiment of the present invention, the router 26 is used to help establish an IPSEC SA by acting as a Local Certificate Authority (“LCA”). In one exemplary preferred embodiment of the present invention, the router 26 acts as an LCA and is itself registered with a higher-level CA. The router 26 itself holds a certificate in which a public encryption key for the router 26 is bound to its global IP 48 address (e.g., IP 48 address 28 (FIG. 1)) that is validated by the higher-level CA. The router 26 acts as a LCA to issue security certificates to other network devices (14, 16, 18, 20, 22, 24) on the first computer network 12 to help establish an IPSEC SA. However, other network devices may also be used as a LCA besides the Router 26.

FIG. 22 is a flow diagram illustrating a Method 304 for providing a security association using distributed NAT. At Step 306, one or more locally unique ports are requested with a first message from a first protocol on a first network device from a second network device. The one or more locally unique ports are used for distributed NAT. At Step 308, one or more locally unique security values are requested with a first message from the first protocol on a first network device from the second network device. The one or more locally unique security values are used with a second secure protocol to establish one or more secure virtual connections between the first network device and a third network device and a second external computer network and for distributed NAT with security. At Step 310, a security certificate is requested

on the first network device from the second network device. The security certificate includes a binding between a public encryption key for the first network device and a combination of a common external network address for the first network device and the one or more locally unique ports allocated by the second network device. The binding comprises a security name space.

In one preferred embodiment of the present invention, the locally unique ports are DNAT ports, the first protocol is the PAP 64, the first message is a PAP 64 security request message 67, and the second secure protocol is IPSEC, and the one or more locally unique security values are SPIs. In one exemplary preferred embodiment of the present invention, IKE may be considered a security protocol within the IPSEC protocol suite. In another embodiment of the present invention, IKE is not considered a security protocol with the IPSEC protocol suite.

IKE is a security protocol that carries a certificate and a SPI value. IKE negotiates a session key that includes a SPI. However, other protocols may also be used to negotiate a session key. The network address is a local IP 48 network address on the first computer network 12 and the second network device is the router 26. However, the present invention is not limited to the ports, protocols, messages, security values, network addresses or network devices discussed, and other ports, protocols, messages, security values, network addresses or network devices could also be used.

In one exemplary preferred embodiment of the present invention, at Step 306, one or more locally unique DNAT ports are requested with a PAP 64 request message 66 on a first network device (e.g., 14) from the router 26 (e.g., with Method 130 of FIG. 9). At Step 308, one or more locally unique SPIs are requested with a PAP 64 security request message 67 from the

Router 26. (e.g., with Method 274 of FIG. 19). The one or more locally unique SPIs are used with IPSEC to establish one or more SAs between the first network device 12 and a third network device 39 and a second external computer network 30. At Step 310, a security certificate is received on the first network device from the router 26. The security certificate includes a binding between the public encryption key and a combination of a common external IP 48 address for the first network device (e.g., 198.10.20.30) and the one or more locally unique DNAT ports allocated to the first network device. The security certificate is used to establish a SA as is described below.

FIG. 23 is a flow diagram illustrating a Method 312 for distributed NAT using security.

A first message with a first protocol from a first network device is received on a second network device to request one or more locally unique ports. The second network device allocates one or more locally unique ports. At Step 314, the second network device sends the allocated one or more locally unique ports to the first network device using a second message from the first protocol. The one or more locally unique ports are used for distributed NAT. A first message with a first protocol from a first network device is received on a second network device to request one or more locally unique security values. The second network device allocates one or more locally unique security values. At Step 316, the second network device sends the allocated one or more locally unique security values to the first network device using a second message from the first protocol. The one or more locally unique security values are used with a second secure protocol to establish a secure virtual connection between the first network device and a third network device and a second external computer network and are used for distributed NAT with security.

A public encryption key and a private encryption key are generated on the first network device. The public encryption key is sent to the second network device from the first network device. The second network device creates a security certificate for the first network device. The security certificate includes a binding between the public encryption key and a combination  
5 of an external network address for the first network device and the one or more locally unique security values. In one exemplary preferred embodiment of the present invention, the security certificate is an Internet X.509 security certificate. However, other types of security certificates could also be used and the present invention is not limited to Internet X.509 security certificates. For more information on Internet X.509 security certificates, see RFC-2459, "Internet X.509  
10 Public Key Infrastructure Certificate and CRL Profile," by R. Housley, W. Ford, W. Polk and D. Solo, incorporated herein by reference. For more information on X.509 security certificate management, see RFC-2510 "Internet X.509 Public Key Infrastructure Certificate Management Protocols," by C. Adams and M. Farrell, and RFC-2511 "Internet X.509 Certificate Request  
15 Message Format", by M. Myer, C. Adams, D. Solo, and D. Kemp, incorporated herein by reference. At Step 318, the second network device sends the security certificate to the first network device.

In one preferred embodiment of the present invention, the locally unique ports are DNAT ports, the first protocol is the PAP 64, the first message is a PAP 64 security request message 67, the second message a PAP 64 security response message 69 the second secure protocol is IPSEC,  
20 the one or more locally unique security values are SPIs, the network address used in the CA is an external IP 48 network address of the second network address on the first computer network 12 and the second network device is the router 26. However, the present invention is not limited to

the ports, protocols, messages, security values, network addresses or network devices discussed, and other ports, protocols, messages, security values, network addresses or network devices could also be used. After receiving one or more locally unique ports, one or more locally unique security values and the security certificate, a network device can use IPSEC with distributed  
5 NAT.

FIG. 24 is a flow diagram illustrating a Method 320 for distributed NAT using security. At Step 322, a first message in a second secure protocol is received on a first network device on a first network including a request to establish a secure connection to the first network device from a third network device on a second external network. At Step 324, a locally unique security  
10 value is selected to use for the secure connection from a stored list of locally unique security values on the first network device. The stored list of locally unique security values was received from a second network device on the first network with a first protocol (e.g., Method 304 of FIG. 22). At Step 326, a second message is sent with the second secure protocol to establish a secure  
15 virtual connection to the first network device on the first network from the third network device on the second external network with the selected locally unique security value and a security certificate received by the first network device. (e.g., at Step 310 of Method 304 (FIG. 22)).

In one preferred embodiment of the present invention, the first network device is a network device (14, 16, 18, 20, 22, and 24) on the first computer network 12. The second network device is the router 26, the third network device is an external network device 39, the  
20 first protocol is the PAP 64, the second protocol is IPSEC, the locally unique security value is a SPI allocated by the router 26 with the PAP 64, and the secure connection is a SA. However, the

present invention is not limited to this exemplary preferred embodiment, and other network devices, protocols, security values and secure connections could also be used with Method 320.

In one exemplary preferred embodiment of the present invention, a network device negotiates an incoming IPSEC SA with a remote network device on an IP 48 network 30. The SPI selected and assigned to a SA is selected from the one or more of locally unique SPI values allocated by a router 26 with PAP 64 to the network device. In one exemplary preferred embodiment of the present invention, an incoming IPSEC SA includes a SA that terminates at the network device for inbound packets (i.e., packets sent from the remote network device to the network device). For outgoing SAs, a SPI is selected by the remote network device and a locally unique SPI is not used by the router 26. In the event of multiple levels of incoming SAs that terminate on a network device, a SPI from the list of locally unique SPI values is allocated only to an outermost SA. A SPI is stored in an IPSEC protocol header of an associated IP 48 packet. For an outermost SA, an IPSEC protocol header is typically visible for combinations of the IPSEC protocol (e.g., AH and ESP) and mode (e.g., transport and tunnel). Thus, the router 26 can access a SPI in an outermost SA associated with any incoming IP 48 packet. After one or more SAs are established between a network device and a remote network device, DNAT with security can be used.

### **Using IPSEC and DNAT**

A first network device on a first network exchanges messages with a third network device on a second external network to establish a security association. For example, the first network device exchanges IKE messages to establish a security association with the external third network device. After exchanging some of these messages, a security value (e.g., SPI) allocated

with PAP 64 will be used to complete the establishment of a security association between the two network devices.

FIG. 25 is a flow diagram illustrating a Method 328 for distributed NAT with security.

At Step 330, a request in a second secure protocol is sent from a first network device on a first network to a second network device on the first network for a third network device on an external second network. The request includes security request information provided to the first network device. In one preferred embodiment of the present invention, the security request information includes a locally unique security value (e.g., SPI) allocated by the second network device with a first protocol (e.g., Method 304 of FIG. 22). The locally unique security value is provided to the first network device by the second network device (e.g., Method 304 of FIG. 22). In another embodiment of the present invention, the security request information includes a security certificate provided by a CA as was discussed above. At Step 332, the request is routed from the second network device to a third network device on a second external network. At Step 334, a response in the second secure protocol is received on the second network device on the first network for the first network device from the third network device on the second external network. The response in the second secure protocol includes security information from the request provided to the first network device. At Step 336, the response is routed from the second network device to the first network device on the first network using a locally unique port from the reply in the second secure protocol. The response completes the establishment of a security association between the first network device and the external third network device using the locally unique security value.



In one preferred embodiment of the present invention, the first network device is a network device (14, 16, 18, 20, 22, and 24) from the first computer network 12, the second network device is the router 26, the first protocol is the PAP 64, the second secure protocol is IPSEC, the locally unique security value is a SPI allocated by the router 26 with the PAP 64, the security association is a SA. In this embodiment of the present invention, IPSEC includes IKE.

As was discussed above, IKE is a protocol that carries a security certificate and a SPI value. IKE negotiates a session key and a SPI that is associated with a session key. However, other protocols can also be used to negotiate a session key. However, the present invention is not limited to this exemplary preferred embodiment, and other network devices, protocols, security values and secure connections could also be used with Method 328.

IKE can be used in two separate modes called the "Main Mode" and "Aggressive Mode." In the Main Mode an SPI is sent in a first and second message (the first from the initiator to the responder, the second from the responder to the initiator) and then security certificates are sent in fifth and sixth messages (the fifth from the initiator to the responder and the sixth from the responder to the initiator). The third and fourth messages are used to continue the IKE negotiations. In the Aggressive mode, on the other hand, the SPI is sent in the first and second messages, while the certificates are sent in the third and fourth messages. The request and response messages in Method 328 can be any of the IKE messages used in the Main mode or the Aggressive mode to send a SPI or a security certificate.

In one exemplary preferred embodiment of the present invention, using IPSEC over DNAT, the router 26 (FIG. 1) does not look at TCP 58 or UDP 60 port numbers for outbound packets, even though they may be visible using IPSEC with AH. For outgoing packets using

IPSEC, the router 26 removes a virtual tunnel header and forwards the remaining packet over an external network interface 28 to an IP 48 network 30. The virtual tunnel header is an outermost header on the data packet.

For incoming packets using IPSEC, the router 26 (FIG. 1) maintains a mapping (FIG. 21) between local IP addresses of network devices (e.g., 14, 16, 18, 20, 22, 24) and SPI values (e.g., Step 288 of Method 282 (FIG. 20). When an AH or ESP IPSEC packet arrives on the router 26, the router 26 examines a SPI value in an IPSEC packet's outermost header. As was discussed above, the outermost IPSEC header is typically visible. The SPI value in the IPSEC header is used to determine a local IP 54 address of a destination network device. A tunneling header is constructed and prepended to the packet (e.g., see Tables 3, 5, 8, and 11). The packet is forwarded to a local network device, and the local network device removes the tunnel header and processes the packet. Thus, the router 26 does not modify contents of a received IPSEC packet.

Even though TCP 58/UDP 60 ports are not used with IPSEC for address mapping by the router 26, they are still used for DNAT once the IPSEC packet is decoded. That is, once IPSEC input processing is complete, DNAT as described above is used (e.g., see FIGS. 9 and 10 and FIGS. 13 and 14 and associated text). Port numbers are also required by a remote second network to properly identify connections to network devices on the first network, in the event that more than one device on the first network has established connections with a remote third network device.

The router 26 is used for both DNAT port and SPI allocation and de-allocation. Local network devices can request additional port numbers and additional SPIs that are allocated by the router 26. The router 26 can also render an allocated range of DNAT ports or SPIs invalid. If

IPSEC is implemented as well, additional security certificates may be issued by the LCA with allocation of additional DNAT ports and SPIs to local network devices. In addition, the router 26 maintains a list of all security certificates issued to its local network devices, and ensures that the associated DNAT ports are never de-allocated as long as the security certificates with  
5 bindings to these DNAT ports are still valid.

Alternatively, if the router 26 is allowed to de-allocate DNAT ports, it revokes any security certificates with bindings to these DNAT ports. Security certificate revocation includes notification to remote systems that have active SAs established with the local network devices whose security certificates have been revoked. De-allocation and security certificate revocation  
10 may be required, for example, when a local network device has a system crash. In the event of a system crash on the router 26, security certificates are sent again to network devices or invalid security certificates are gracefully revoked.

The methods of authentication are not restricted to the form of the name space for binding of security certificates described above. For example, a combination of the router's 26  
15 global IP 48 address 28 and a user e-mail address (where the user is on a local network device) could also be used for a name space binding for a security certificate. The router 26 acting as an LCA should possess a valid security certificate giving it the right to certify identifiers drawn from a chosen name space.

The methods for preferred embodiments of the present invention presented herein also  
20 extends IPSEC within the context of Mobile IP, allowing a mobile node to maintain an IPSEC-protected connection while it is roaming. For Mobile IP, a mobile node's home agent's global IP address and a mobile node's local address on its home network can be used for name space

binding to create a security certificate to use for IPSEC with DNAT. This information is available to a mobile node even while it is roaming (i.e., temporarily residing on a foreign network). A mobile node's home network is managed as a DNAT stub network in which the mobile node resides as a local host when it is not roaming. Using DNAT with Mobile IP is described in co-pending Application No. 09/136,484.

A modified security name space can be used to provide a unique identifier in a security certificate to a network device that lacks a globally unique IP 48 address and is not restricted to a design based upon the router 26 acting as an LCA. It also is possible to define a global CA using a modified name space, and eliminate the need for the LCA, or the router 26 acting as a LCA.

However, such a modified name space is typically insufficient for the DNAT environment, since it does not include a locally unique port number, and hence does not guarantee to a remote system that a local network device has the right to use a specific port number. Also, since stub networks exist, and DNAT includes methods for sharing global IP 48 addresses within stub networks, the LCA approach described herein provides an implementation that would build upon an existing infrastructure, rather than requiring a new infrastructure if a DNAT system is used. Thus, IPSEC can be used with DNAT with the router 26 acting as an LCA without requiring a new infrastructure to support a global CA.

### **Controlling Denial Of Service Attacks With DNAT and IPSEC**

As set forth above, a first network device (e.g., 14, FIGS. 1 and 26) on a first computer network 12 may use a second network device on the first computer network, such as DNAT router 26, for secure communication (i.e., an IPSEC SA) with a third external network device (e.g., 39, FIGS. 1 and 26) on a second external computer network 30 (e.g., the Internet/Intranet,

FIGS. 1 and 26). Through PAP 64, the second network device (i.e., DNAT router 26) will then allocate a range of TCP/UDP port numbers and security values, such as IPSEC SPIs, for the first network device to utilize. Incoming packets to the second network device (i.e., DNAT router 26) that have either an ESP or an AH header following an IP Header will be sent to the proper network device on the first computer network (e.g., the first network device) by looking up the SPI of the incoming packet in an SPI-to-internal network address table (e.g., layout 292, FIG. 21) and sending the packet to the network device on the first computer network to which that SPI is allocated.

As a result, however, the first network device and the DNAT-enabled first computer network 12 may be susceptible to denial of service attacks from a fourth external network device 41 that is connected to the second external computer network 30, as shown in FIG. 26, and with which no SA has been established. By the fourth external network device 41 transmitting IPSEC packets using an SPI that belongs to the first network device and an IP address that belongs to the second network device and is shared with the first network device, these packets will be forwarded by the second network device (i.e., DNAT router 26) to the first network device. The first network device's IPSEC implementation will preferably discard these packets. However, the fourth external network device 41 may transmit hundreds or thousands of packets in rapid succession, thereby swamping resources in one or more of the following locations: (1) the second network device (i.e., DNAT router 26), (2) the first computer network 12, and (3) the first network device (e.g., 14, FIGS. 1 and 26). This swamping of resources is well-known in the art as a Denial of Service (DoS) attack.

DNAT may be used with IPSEC, however, to minimize, control, and limit the disruption from DoS attacks caused by a flood of packets from external network devices. A preferred method of the present invention for doing so is set forth in the methods 400, 500, which are illustrated by the block diagrams in FIGS. 27 and 29. It should be understood that the methods 400, 500 of the present invention together form the preferred method of the present invention for controlling and limiting the disruption from DoS attacks, and that the methods 400, 500 are shown and described below as separate methods for ease of explanation and illustration only.

In Step 410 of the method 400, the first network device (e.g., 14, FIG. 26) on the first computer network 12 establishes an SA with a third external network device (e.g., third external network device 39, FIG. 26) on the second external computer network 30 using DNAT and IPSEC as previously described. Next, in Step 412, the first network device specifies to a second network device (e.g., DNAT router 26) on the first computer network 12, the external IP addresses (e.g., 192.200.20.3, FIG. 26) of one or more external network devices with which the first network device has established SAs. These external IP addresses for the established SAs are considered to be valid by the first and second network devices on the first computer network. As explained in more detail below, an extension of PAP 64 may be used by the first network device to specify the external IP addresses for the established SAs to the second network device.

The second network device then stores the external IP addresses in a table for established SAs in Step 414. There are at least three possible strategies for such a table. The second network device (i.e., DNAT router 26) can: (1) associate external IP addresses with internal network devices on the first computer network, (2) associate external IP addresses with security values, such as SPIs, or (3) associate external IP addresses with both internal network devices

and security values, such as SPIs. After the second network device (i.e., DNAT router 26) enters the external IP address into the table, the second network device maps the external IP address in the table to the internal network address of the first network device and/or to the security value (i.e., SPI) for the established SA in Step 416. Examples of suitable tables for use with Steps 414 and 416 are shown in FIGS. 29A-C and described in more detail below.

In one exemplary preferred embodiment of the present invention, the valid external IP address is specified and delivered by the first network device to the second network device using an extension of PAP 64 having a PAP external address validating message 420. Preferably, the PAP external address validating message 420 is sent immediately upon the establishment of an SA. A preferred layout for this validating message 420 is shown in FIG. 28A. The validating message 420 has a type-field 421 that is one-byte and has a value (e.g., 34) for specifying a valid external IP address associated with an established SA. The validating message 420 also has a code-field 422 that is one-byte and has a value of zero for the valid external IP address, and a checksum-field 423 that is two-bytes and has a value of a 1's complement sum of the entire PAP external address validating message 420 layout. Preferably, the validating message 420 further includes a valid external address field 424 that is four-bytes and has a variable value indicating a valid external IP address for which an SA has been established.

As shown in FIG. 28A, the validating message 420 may also include a valid SPI field 425 that is four-bytes and has a variable value indicating a valid SPI associated with the established SA. It should be understood, however, that other layouts, values and field sizes could also be used for the PAP external address validating message 420 layout.

The first network device uses the PAP external address validating message 420 to notify the second network device (e.g., DNAT router 26) of which external IP addresses the first network device has established an SA with. The first network device may also be able to invalidate this assignment of external IP addresses, and disassociate any specified SPIs from their corresponding network addresses. In one exemplary preferred embodiment of the present invention, the external IP address tabled by the second network device (e.g., DNAT router 26) may be invalidated by the first network device using an extension of PAP 64 having a PAP external address invalidating message 430. Preferably, the PAP external address invalidating message 430 is sent immediately upon the termination of an SA.

A preferred layout for this invalidating message 430 is shown in FIG. 28B. The invalidating message 430 has a type-field 431 that is one-byte and has a value (e.g., 35) for specifying an invalid external IP address that is no longer associated with an SA. The invalidating message 430 also has a code-field 432 that is one-byte and has a value of zero for the invalid external IP address, and a checksum-field 433 that is two-bytes and has a value of a 1's complement sum of the entire PAP external address invalidating message 430 layout. Preferably, the invalidating message 430 further includes a invalid external address field 434 that is four-bytes and has a variable value indicating an invalid external IP address for which an SA has been terminated. As shown in FIG. 28B, the invalidating message 430 may also include an invalid SPI field 435 that is four-bytes and has a variable value indicating a invalid SPI that is no longer associated with an established SA. It should be understood, however, that other layouts, values and field sizes could also be used for the PAP external address invalidating message 430 layout.



As noted above, examples of suitable established SA tables for use with Steps 414 and 416 are shown in FIGS. 29A-C. It should be understood that while each of the established SA tables may have multiple rows, only the first row for each table is shown in FIGS. 29A-C for ease of illustration. In FIG. 29A, the external IP addresses of any external network devices for which an SA has been established with an internal network device may be stored in an external network address for SA column 452a of an established SA table 450a. For example, the external IP address of the third external network device 39 (e.g., 192.200.20.3) is stored in the external network address for SA column 452a and in the first row 456a of the established SA table 450a, as shown in FIG. 29A.

The internal network address of the internal network device that has established an SA with the tabled external network device may be stored in an internal network address column 454a of the established SA table 450a, and mapped to the corresponding external IP address for the SA, as shown in FIG. 29A. For example, assuming that an SA has been established between the first network device (e.g., 14, FIG. 26) and the third external network device (e.g., 39, FIG. 26), the internal network address (e.g., 10.0.0.1) of the first network device would be stored in the internal network address column 454a and in the first row 456a of the established SA table 450a, and mapped to the external IP address (e.g., 192.200.20.3) of the third external network device stored in the external network address for SA column 452a and in the first row 456a of the established SA table 450a.

FIGS. 29B and 29C illustrate alternative examples for established SA tables. In FIG. 29B, the established SA table 450b is identical to the established SA table 450a, except that the established SA table 450b has an SPI for SA column 452b in place of the internal network

address column 454a. As a result, while the external network address for SA column 452b is identical to the external network address for SA column 452a, the first row 456b of established SA table 450b includes the SPI for the SA established between the internal and external network devices, in addition to the corresponding external IP address of the third external network device.

5 For example, assuming that an SA has been established between the first network device (e.g., 14, FIG. 26) and the third external network device (e.g., 39, FIG. 26) using an SPI of 280, the external IP address (e.g., 192.200.20.3) of the third external network device would be stored in the external network address for SA column 452b and in the first row 456b of the established SA table 450b, and mapped to the SPI (e.g., 280) for the SA stored in the SPI for SA column 454b and in the first row 456b of the established SA table 450b. It should be understood that the SPI stored in the SPI for SA column 454b may be further associated with the internal network address for the internal network device of the SA through the SPI-to-internal network address table (e.g., layout 292, FIG. 21) described above.

The established SA table 450c shown in FIG. 29C is a combination of the established SA tables 450a, 450b shown in FIGS. 29A-B. The external network address for SA column 452c is identical to the external network address for SA columns 452a, 452b, the SPI for SA column 454c is identical to the SPI for SA column 454b, and the internal network address column 458c is identical to the internal network address column 454a. As a result, the first row 456c of the established SA table 450c includes the external IP address (e.g., 192.200.20.3) of the external network device (e.g., 39, FIG. 26) of the SA, the SPI (e.g., 280) for the SA, and the internal network address (e.g., 10.0.0.1) of the internal network device (e.g., 14, FIG. 26) of the SA.

5 The method 500 of the present invention shown in FIG. 30 further illustrates how DNAT and IPSEC may be used to control and limit the disruption from DoS attacks caused by a flood of packets from external network devices. In Step 510, an external network device (e.g., 39 or 41, FIG. 26) on a second computer network (e.g., 30, FIG. 26) sends a data packet to an internal network device (e.g., 14, FIG. 26) on a first computer network (e.g., 12, FIG. 26) using the above described DNAT with security (i.e, IPSEC). Since DNAT is being implemented, the packet is intercepted by the DNAT router (e.g., 26, FIG. 26) on the first computer network in Step 520. In Step 520, the DNAT router also determines the external IP source address and SPI of the packet. The DNAT then determines whether the SPI has been allocated to the internal network device in Step 530 by looking up the SPI in the SPI-to-internal network address table (e.g., layout 292, FIG. 21). If the SPI has not been allocated to the internal network device, the DNAT router discards the packet in Step 540.

10 If the SPI has been allocated to the internal network device, the DNAT router determines in Step 550 whether the external IP address of the packet is valid, i.e., the external IP address of the packet has been specified by the internal network device as being associated with an established SA. If the external IP address of the packet is not valid, i.e., the external address of the packet was not specified by the internal network device as being associated with an established SA, then the DNAT router discards the packet in Step 540. On the other hand, if the external IP address of the packet is valid, then the packet is sent by the DNAT router in Step 560 to the appropriate internal network device using the DNAT and IPSEC methods described above.

20 To determine whether the external IP address of the packet is valid in Step 550, the DNAT router preferably looks up the external IP address of the packet on the established SA

table used by the DNAT router. For example, the DNAT router may look up the external IP address of the packet on one of the established SA tables 450a, 450b, 450c to determine whether the external IP address is mapped to the internal network address of the internal network device and/or the SPI for the established SA that has been allocated to the internal network device. If the destination of the packet is the internal network device mapped to the external IP address in the established SA tables 450a, 450c, or the SPI of the packet is mapped to the external IP address in the established SA tables 450b, 450c, then the external IP address of the packet is considered valid, and the packet is forwarded to the internal network device in Step 560 using DNAT and IPSEC.

To help further illustrate the method 500, the flow of packets being sent to the first network device 14 from the third external network device 39, for which an SA has been established, and from the fourth network device 41, for which no SA has been established, will now be described with reference to FIGS. 26, 29A-C, and 30. A packet sent from the third external network device 39 in Step 510 will be intercepted by DNAT router 26 in Step 520. The DNAT router 26 will also then determine the external IP source address (e.g., 192.200.20.3) and SPI (e.g., 280) of the packet from the third external network device 39 in Step 520. Since the SPI of 280 has been allocated to the first network device 14, as shown in the SPI-to-internal network address table (e.g., layout 292, FIG. 21), the DNAT router 26 moves from Step 530 to Step 550. The DNAT router then determines whether the external IP source address of 192.200.20.3 is valid in Step 550 by looking up the external IP source address (or the SPI of 280) on the established SA table 450a (or the other established SA tables 450b, 450c). Because the external IP source address of 192.200.20.3 is mapped in the established SA table 450a (and the

established SA table 450c) to the internal address (e.g., 10.0.0.1) of the first network device 14 (and is also mapped to the SPI of 280 in the established SA table 450b), the external IP source address is considered valid, and the packet is forwarded in Step 560 to the first network device 14 by the DNAT router 26.

5 In contrast, a packet sent from the fourth external network device 41 in Step 510 will be intercepted by DNAT router 26 in Step 520. The DNAT router 26 will also then determine the external IP source address (e.g., 190.100.10.1) and SPI (e.g., 290) of the packet from the fourth external network device 39 in Step 520. Since the SPI of 290 has been allocated to the first network device 14, as shown in the SPI-to-internal network address table (e.g., layout 292, FIG. 10 21), the DNAT router 26 will move from Step 530 to Step 550. The DNAT router then determines whether the external IP source address of 190.100.10.1 is valid in Step 550 by looking up the external IP source address (or the SPI of 290) on the established SA table 450a (or the other established SA tables 450b, 450c). Because the external IP source address of 190.100.10.1 is not mapped in the established SA table 450a (or the established SA table 450c) 15 to the internal address (e.g., 10.0.0.1) of the first network device 14, however, the external IP source address is not considered valid, and the DNAT router discards the packet in Step 540. In addition, although the SPI of 290 was allocated to the first network device 14, the SPI of 290 was not listed by the first network device 14 in the established SA table 450b of the DNAT router 26 as being valid (i.e., mapped to the external IP source address of the fourth external network 20 device 41, 190.100.10.1). Accordingly, the DNAT router would discard the packet in Step 540 based on this SPI lookup (Step 550) as well.

While the preferred method of the present invention (i.e., 400, 500) may not prevent the DNAT router 26 (i.e., the second network device on the first computer network) from being the target of a DoS attack, the method of the present invention does prevent the internal network devices (i.e., the first network device on the first computer network) from being directly affected by such an attack. In other words, the method of the present invention may be used to limit and control the potential disruption caused by a DoS attack. For instance, if the DNAT router 26 is swamped and overloaded with packets from a non-secure external network device (e.g., fourth external network device 41), the DNAT router 26 may not be able to route the internal network devices (e.g., first network device 14) to external computer networks (e.g., 30 and 32, FIG. 26), but the internal network devices (e.g., 14, 16, 18, 20, 22, and 24, FIG. 26) will still be able to communicate internally with each other.

It should be understood that in order to protect against DoS attacks from external network devices that spoof the external source IP address of another external network device that has established an SA with an internal network device, the preferred method of the present invention may be modified to allow the internal network device to inform the DNAT router of external IP addresses that the internal network device does not want to receive packets from, i.e., which external IP addresses are invalid. This may be done in the same manner as the internal network device specifies to the DNAT router which external IP addresses are valid, except that invalid (rather than valid) external IP addresses and/or SPIs would be specified by the internal network device. By letting the internal network device explicitly blacklist external IP addresses and/or SPIs, the internal network device is capable of making a decision of what to do against this type

of DoS attack (i.e., stop traffic that is spoofed, thereby stopping legitimate traffic, or choosing not to stop it, thereby continuing to be a DoS target, but also allowing legitimate traffic through).

It should also be understood that the method of the present invention may be able to provide some protection against a malicious internal network device on a first computer network from colluding with an external network device on a second computer network to perform a DoS attack on the DNAT router and other internal network devices on the first computer network.

Upon detection of such an attack, the DNAT router used with the method of the present invention can easily revoke all resources (e.g., SPIs and/or ports) allocated to the malicious internal network device, and refuse to serve the malicious internal network device further. As a result, the DoS attack will be limited solely to the DNAT router from the external network device on the second computer network, rather than the other internal network devices on the first computer network.

It should be understood that the programs, processes, methods and systems described herein are not related or limited to any particular type of computer or network system (hardware or software), unless indicated otherwise. Various types of general purpose or specialized computer systems may be used with or perform operations in accordance with the teachings described herein.

In view of the wide variety of embodiments to which the principles of the present invention can be applied, it should be understood that the illustrated embodiments are exemplary only, and should not be taken as limiting the scope of the present invention. For example, the steps of the flow diagrams may be taken in sequences other than those described, and more or fewer elements may be used in the block diagrams. In addition, the PAP external address

invalidating message shown in FIG. 28B may also be sent to the DNAT router by a management station or other device on the first computer network. While various elements of the preferred embodiments have been described as being implemented in software, in other embodiments in hardware or firmware implementations may alternatively be used, and vice-versa.

- 5           The claims should not be read as limited to the described order or elements unless stated to that effect. Therefore, all embodiments that come within the scope and spirit of the following claims and equivalents thereto are claimed as the invention.



**WE CLAIM:**

1. A method for distributed network address translation with security, comprising the  
5 following steps:

providing a first network device and a second network device on a first network;

establishing a security association between the first network device and a third network  
device on a second network external to the first network;

specifying an external address of the third network device for the security association;

10 storing the external address in a table on the second network device; and

mapping at least one of an internal address and a security value to the external address in  
the table.

2. A computer readable medium having stored therein instructions for causing a central  
15 processing unit to execute the method of claim 1.

3. The method of claim 1 wherein the second network device is a distributed network  
address translation router.

20 4. The method of claim 1 wherein the security value is a security parameter index for an  
Internet Protocol security protocol.

5. The method of claim 4 wherein the Internet Protocol security protocol is any of an Authentication Header protocol, Encapsulated Security Payload protocol, or an Internet Key Exchange protocol.

6. The method of claim 1 further comprising the step of specifying the external address of the third network device for the security association with a Port Allocation Protocol external address validating message sent from the first network device to the second network device.

7. The method of claim 6 wherein the Port Allocation Protocol external address validating message has a valid external address field.

8. The method of claim 1 further comprising the step of removing the external address from the table with a Port Allocation Protocol external address invalidating message sent from the first network device to the second network device.

9. The method of claim 8 wherein the Port Allocation Protocol external address invalidating message has an invalid external address field.

10. A method for distributed network address translation with security, comprising the following steps:

providing a first network device and a second network device on a first network, and a third network device on a second network external to the first network;

6  
8  
sending a packet having an external address and a security value from the third network  
device to the first network device;

intercepting the packet with the second network device;

determining whether the security value of the packet has been allocated to the first  
5 network device;

determining whether the external address of the packet has been specified by the first  
network device as being valid; and

10 sending the packet from the second network device to the first network device if the  
security value has been allocated to the first network device and the external address of the  
packet has been specified by the first network device as valid.

11. A computer readable medium having stored therein instructions for causing a central  
processing unit to execute the method of claim 10.

15 12. The method of claim 10 wherein the second network device is a distributed network  
address translation router.

13. The method of claim 10 wherein the security value is a security parameter index for  
an Internet Protocol security protocol.

20 14. The method of claim 13 wherein the Internet Protocol security protocol is either an  
Authentication Header protocol or an Encapsulated Security Payload protocol.

15. The method of claim 10 further comprising the step of discarding the packet if the security value of the packet has not been allocated to the first network device.

5 16. The method of claim 10 further comprising the step of discarding the packet if the external address of the packet has not been specified by the first network device as being valid.

10 17. The method of claim 10 further comprising the steps of discarding the packet if the security value of the packet has not been allocated to the first network device, and discarding the packet if the external address of the packet has not been specified by the first network device as being valid.

15 18. The method of claim 10 further comprising the step of specifying the external address as being valid if a security association has been established between the first network device and the third network device.

19. The method of claim 18 further comprising the step of storing a valid external address in a table on the second network device.

20 20. A system for distributed network address translation with security comprising:  
a routing network device using distributed network address translation with security to provide routing services for a plurality of internal and external network devices; and


an established security association table associated with the routing network device for storing external addresses of external network devices that have established security associations with internal network devices, and mapping external addresses that have been specified as valid by the internal network devices to one of internal network addresses and security values for

5 established security associations.

## ABSTRACT OF THE DISCLOSURE

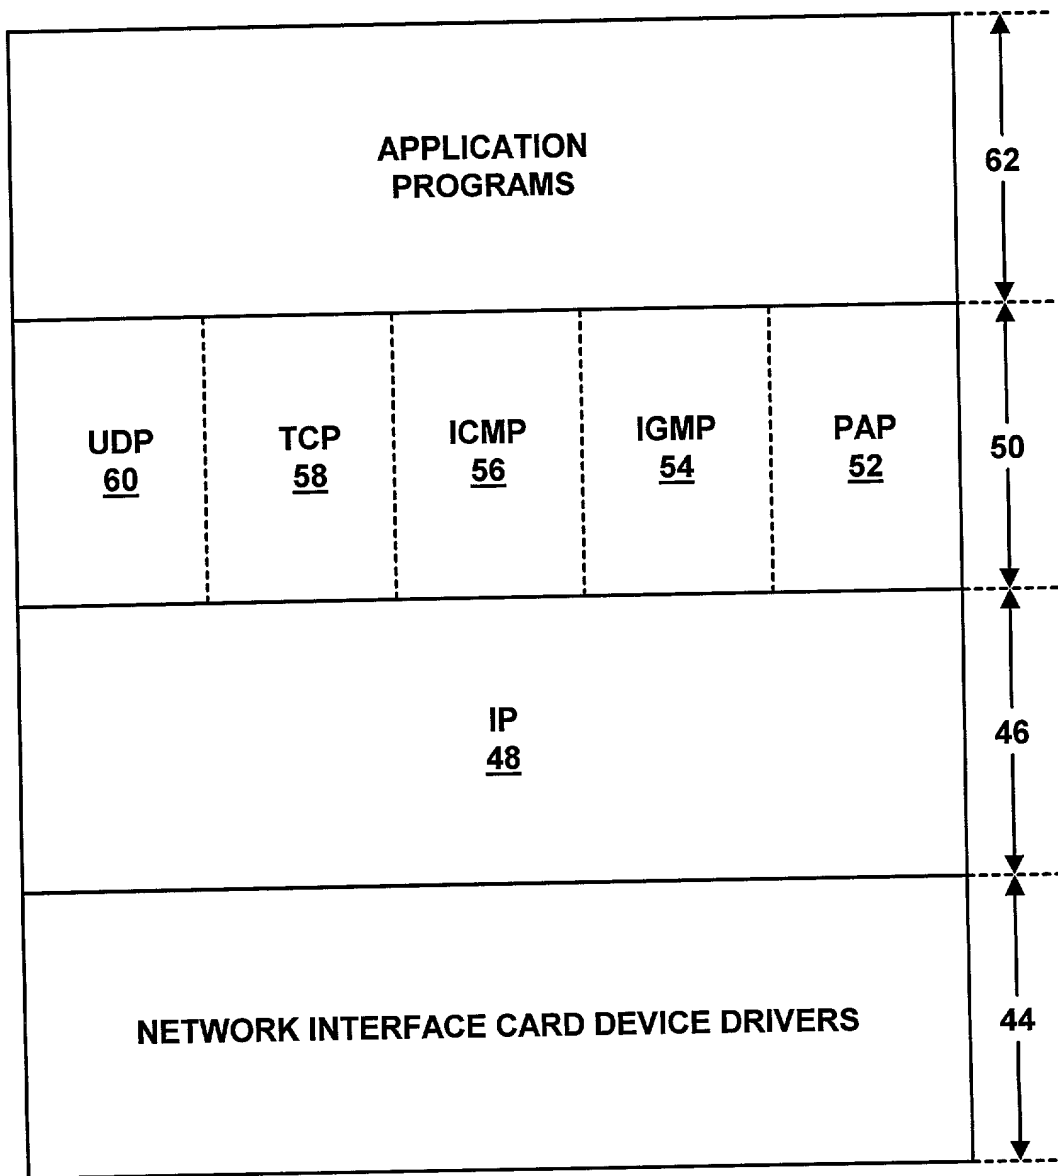
A method and system for distributed network address translation with security for controlling and limiting the disruption caused by denial of service attacks. The method and system have a first network device and a second network device on a first network, and a third network device on a second network external to the first network, with an established security association between the first network device and the third network device. The first network device specifies an external address of the third network device for the security association to the second network device, which stores the external address in a table. The second network device then maps at least one of an internal address and a security value to the external address in the table. Any packets sent from the third network device to the first network device are intercepted by the second network device, which determines the external address and security value of the packet. If the security value of the packet has been allocated to the first network device, and the external address of the packet has been specified by the first network device as being valid, the packet is sent from the second network device to the first network device using distributed network address translation with security. Otherwise, the packet is discarded by the second network device.

FIG. 1



**FIG. 2**  
**PROTOCOL STACK**

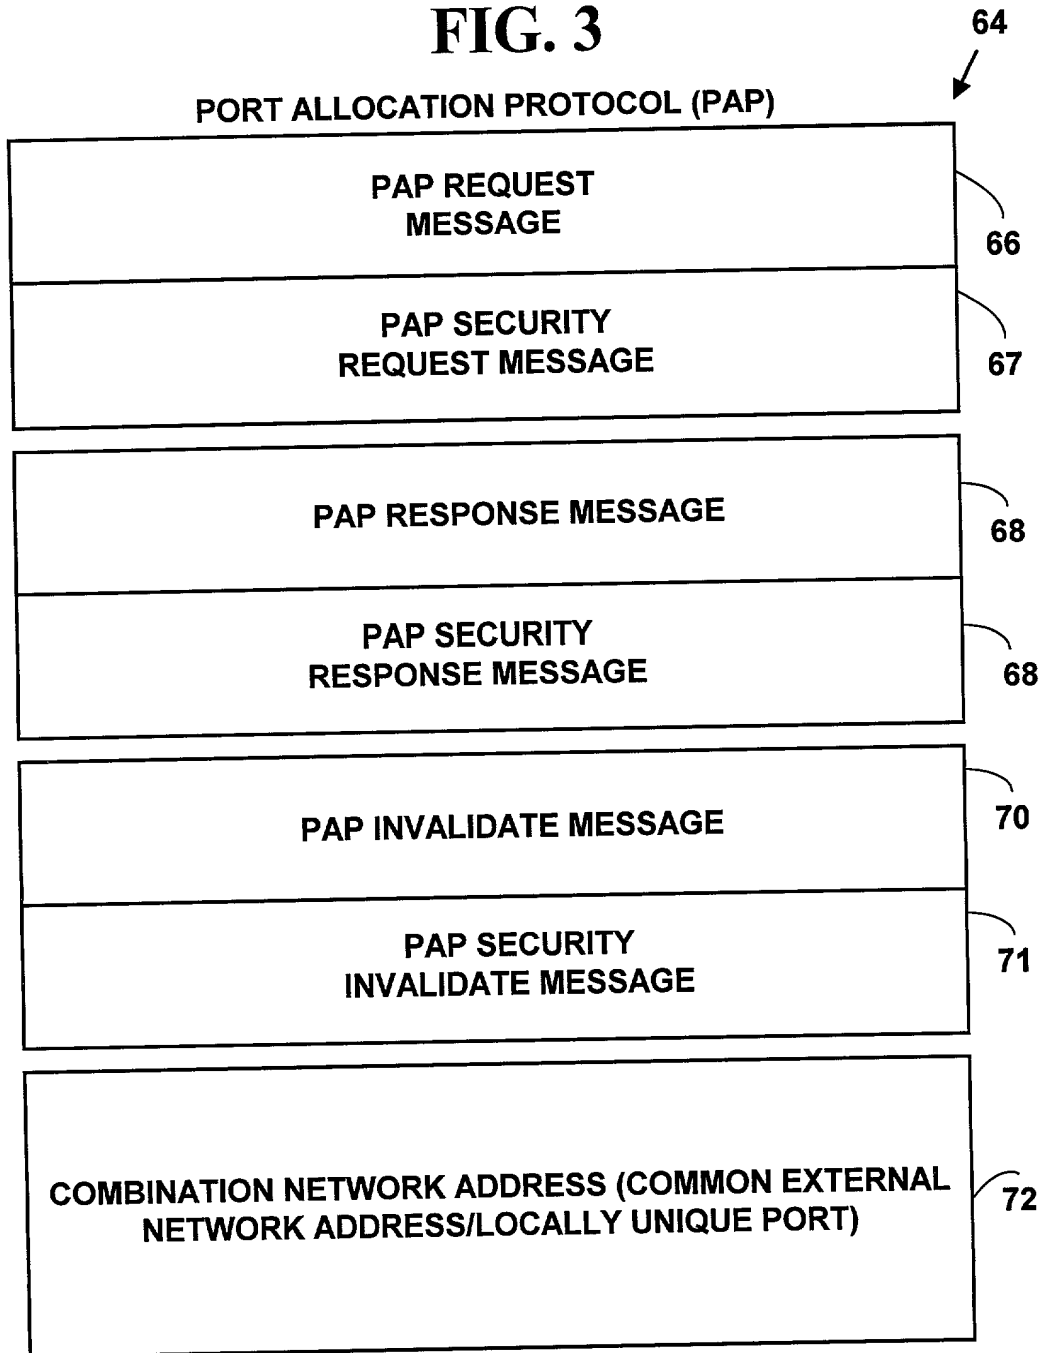
42






**FIG. 3**


**PORT ALLOCATION PROTOCOL (PAP)**




**FIG. 4**



**FIG. 5**



**FIG. 6**



**FIG. 7**

112

114 COMBINATION NETWORK ADDRESS 116

EXTERNAL NETWORK ADDRESS (E.G., EXTERNAL IP ADDRESS)	LOCALLY UNIQUE PORT
198.10.20.30	1032

**FIG. 8**

118


120 122 124

INTERNAL NETWORK ADDRESS	LOWEST PORT	NUMBER OF PORTS
10.0.0.1	1026	32
10.0.0.3	1058	16


126 128

**PORT-TO-INTERNAL-NETWORK  
ADDRESS TABLE**


**FIG. 4A**



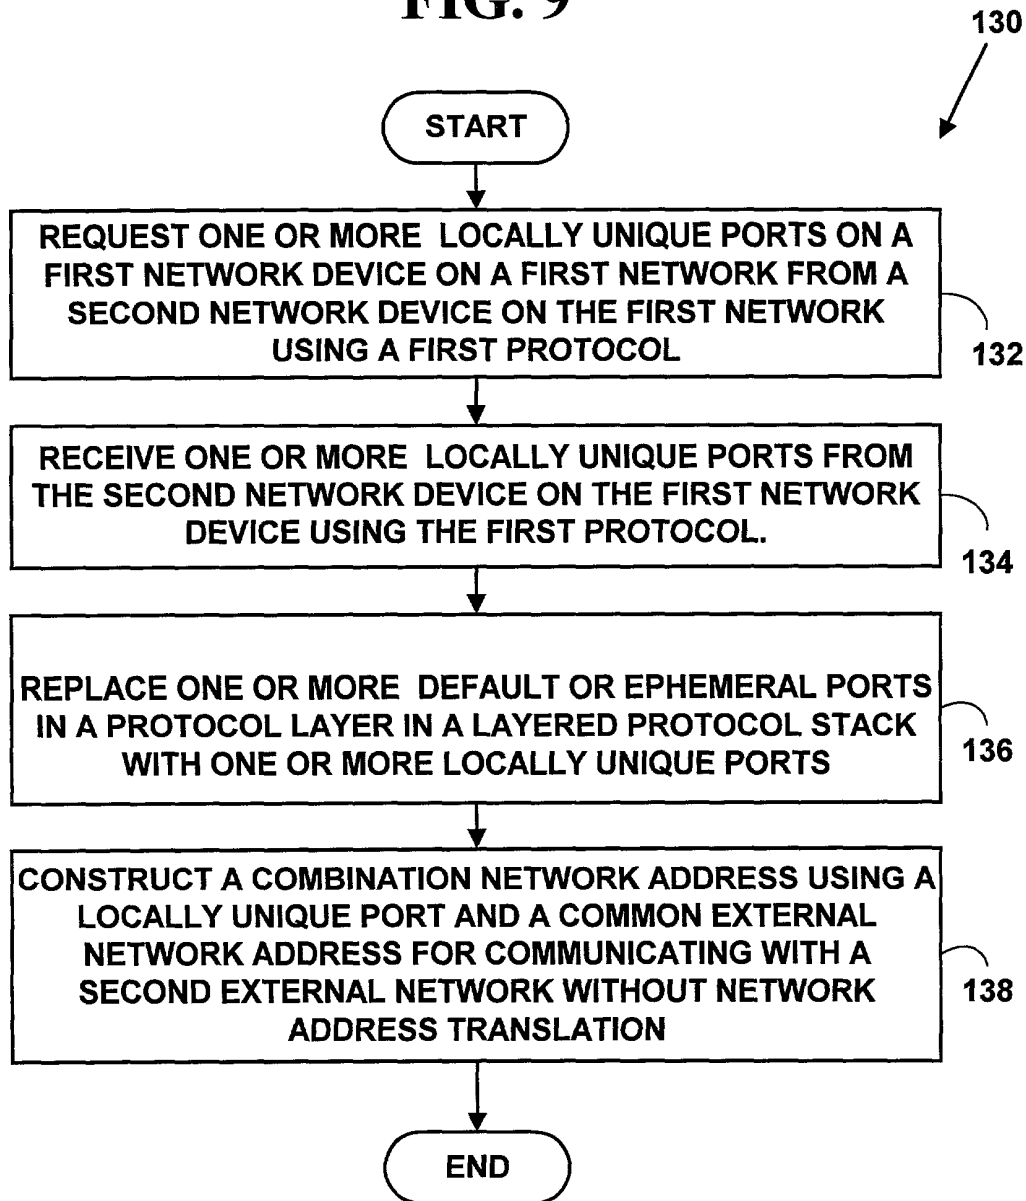
**FIG. 5A**



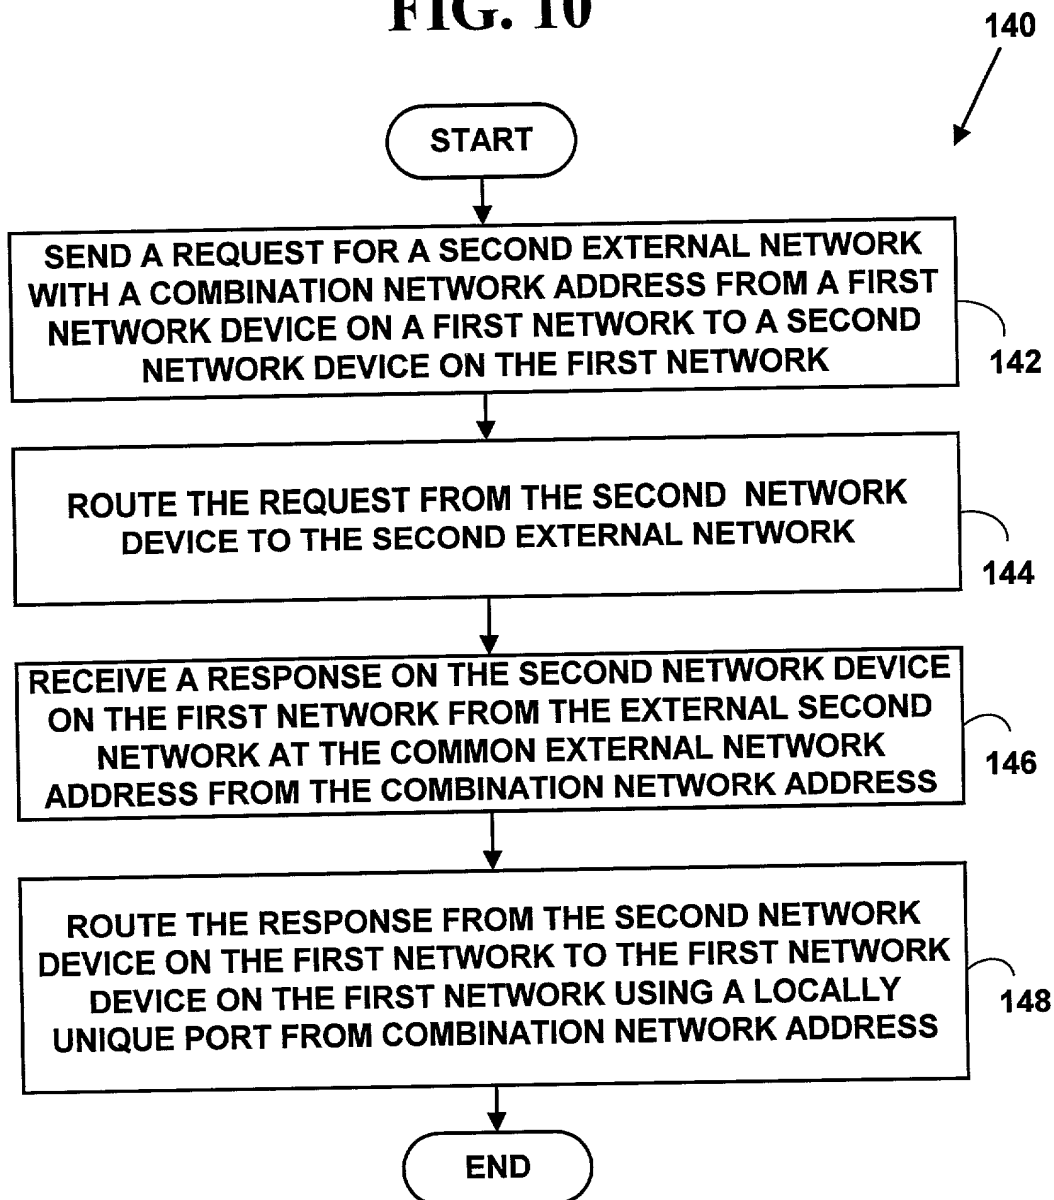
**FIG. 6A**



**FIG. 9**



**FIG. 10**



**FIG. 11**  
SOURCE PORT TRANSLATION TABLE  
(SPTT)

150

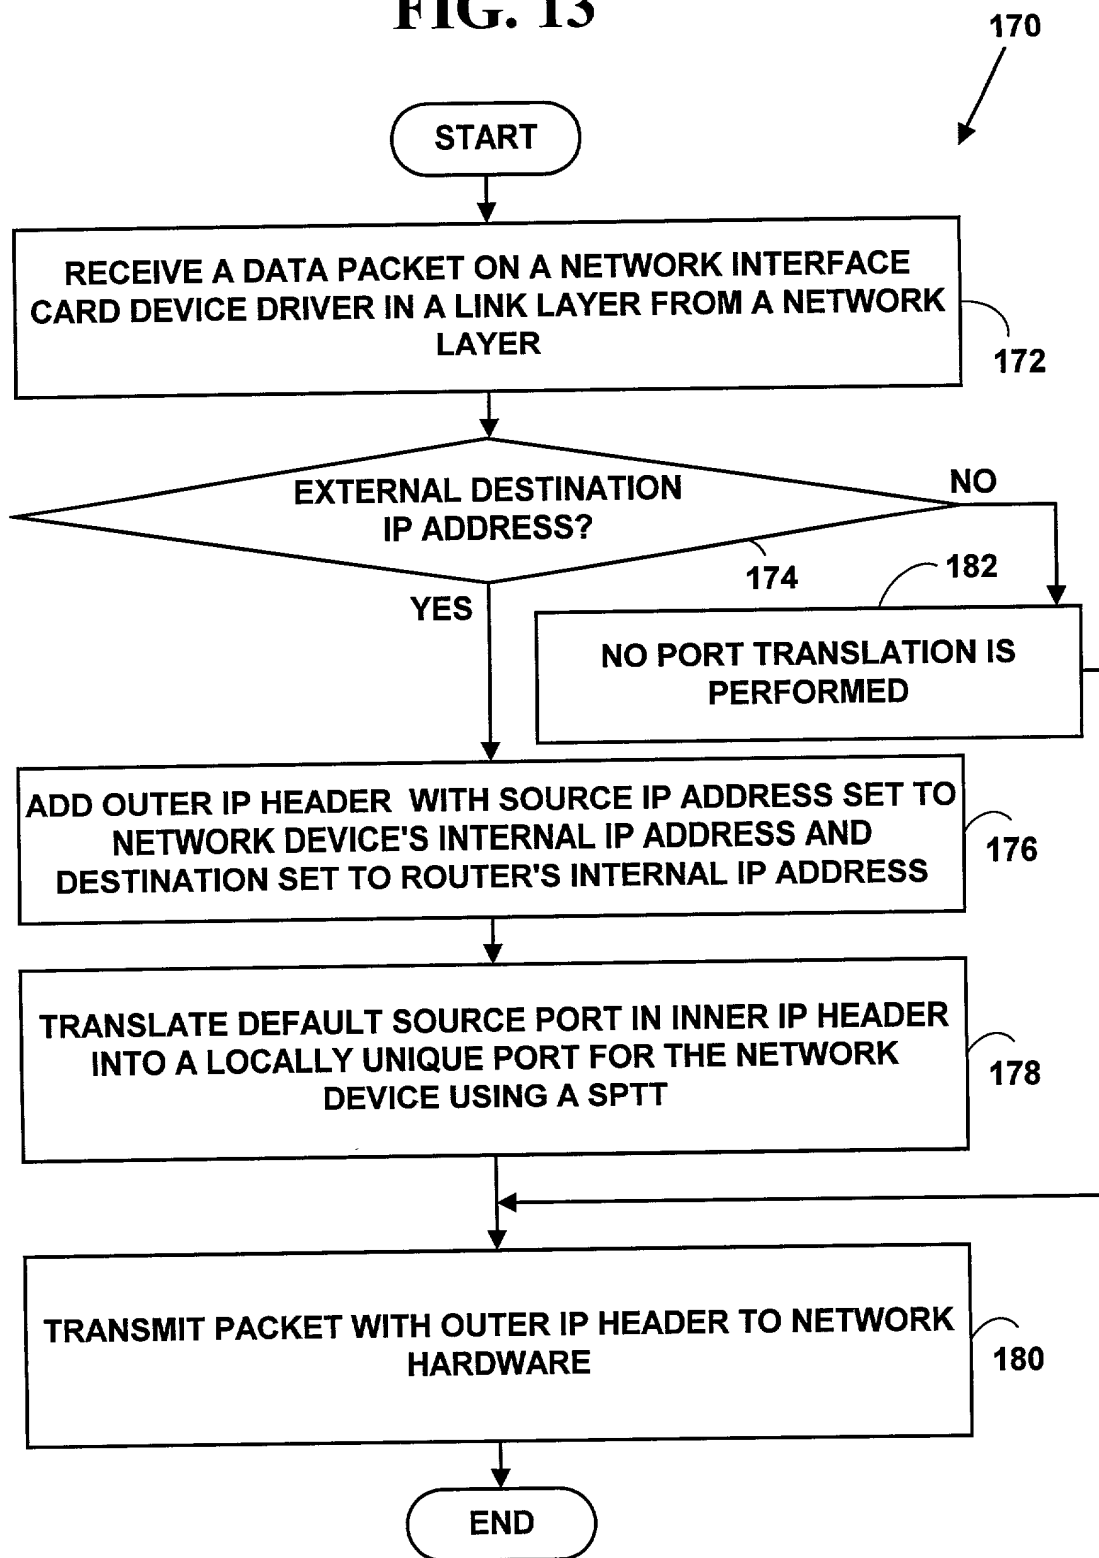
152		154	
DEFAULT PORT		TRANSLATED PORT	
156	PROTOCOL	158	
		TIMESTAMP	

**FIG. 12**  
IP ADDRESS TRANSLATION TABLE  
(IPATT)

160

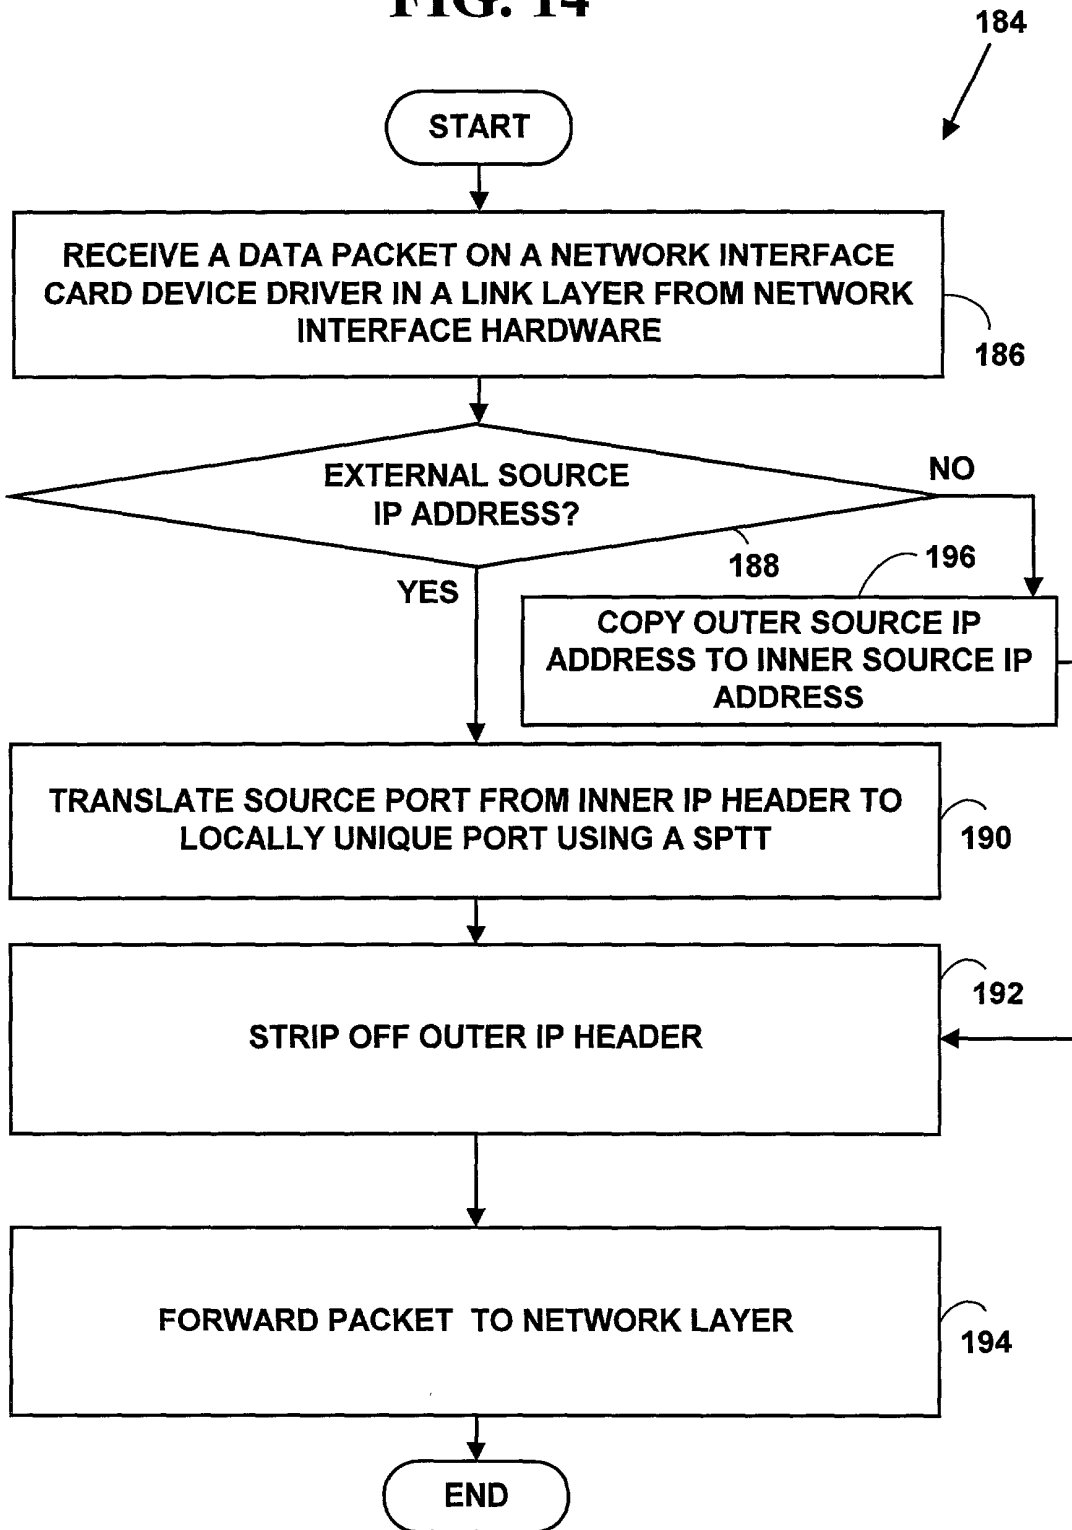
162		164	
DESTINATION PORT		INTERNAL DESTINATION IP ADDRESS	
166	PROTOCOL	168	
		TIMESTAMP	

**FIG. 13**





**FIG. 14**



# FIG. 15

## IP HEADER FORMAT

200



VER <u>202</u>	IHL <u>204</u>	TOS <u>206</u>	TOTAL LENGTH <u>208</u>
IDENTIFICATION <u>210</u>			FRAGMENT OFFSET <u>212</u>
TIME-TO-LIVE <u>214</u>	PROTOCOL <u>216</u>		HEADER CHECKSUM <u>218</u>
SOURCE ADDRESS <u>220</u>			
DESTINATION ADDRESS <u>222</u>			
OPTIONS <u>224</u>			

**FIG. 16**

**AH HEADER FORMAT**

**226**

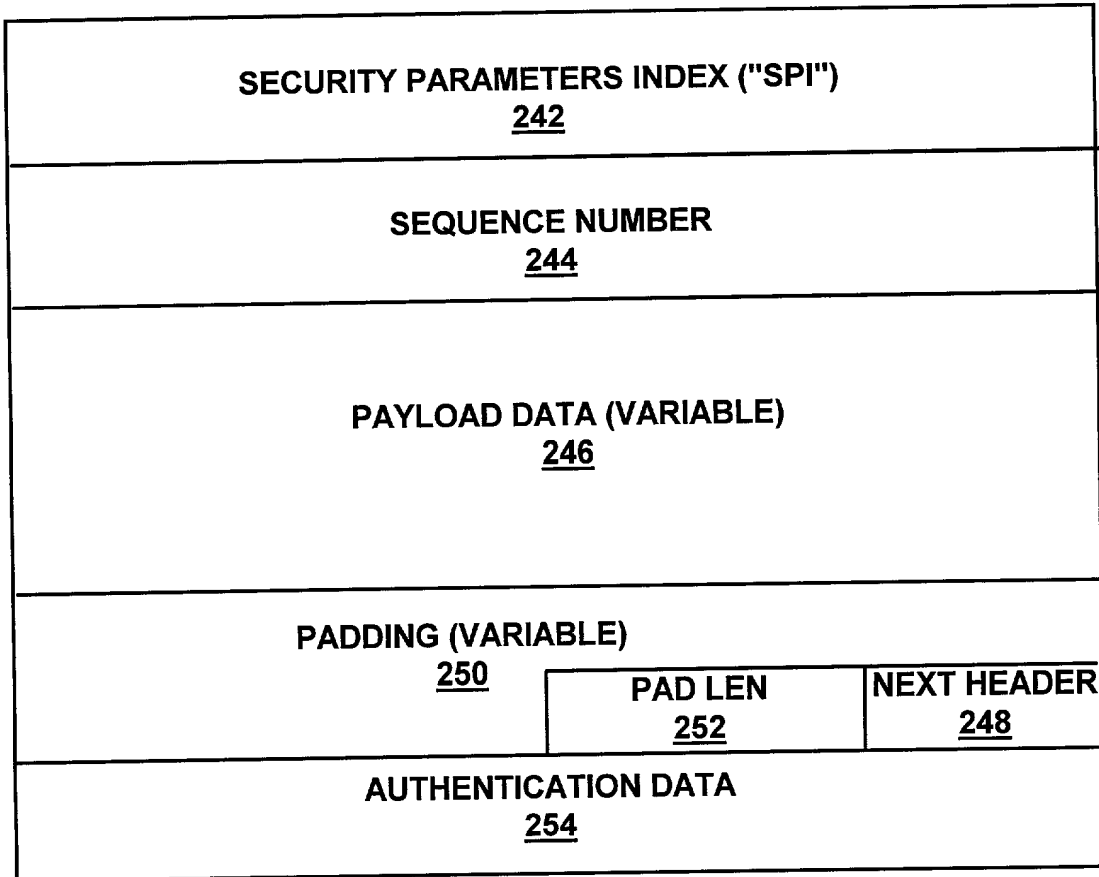


<b>NEXT HEADER <u>228</u></b>	<b>PAYLOAD LENGTH <u>230</u></b>	<b>RESERVED <u>232</u></b>
<b>SECURITY PARAMETER INDEX ("SPI") <u>234</u></b>		
<b>SEQUENCE NUMBER FIELD <u>236</u></b>		
<b>AUTHENTICATION DATA <u>238</u></b>		

**FIG. 17**


**ESP HEADER FORMAT**

240  
↙

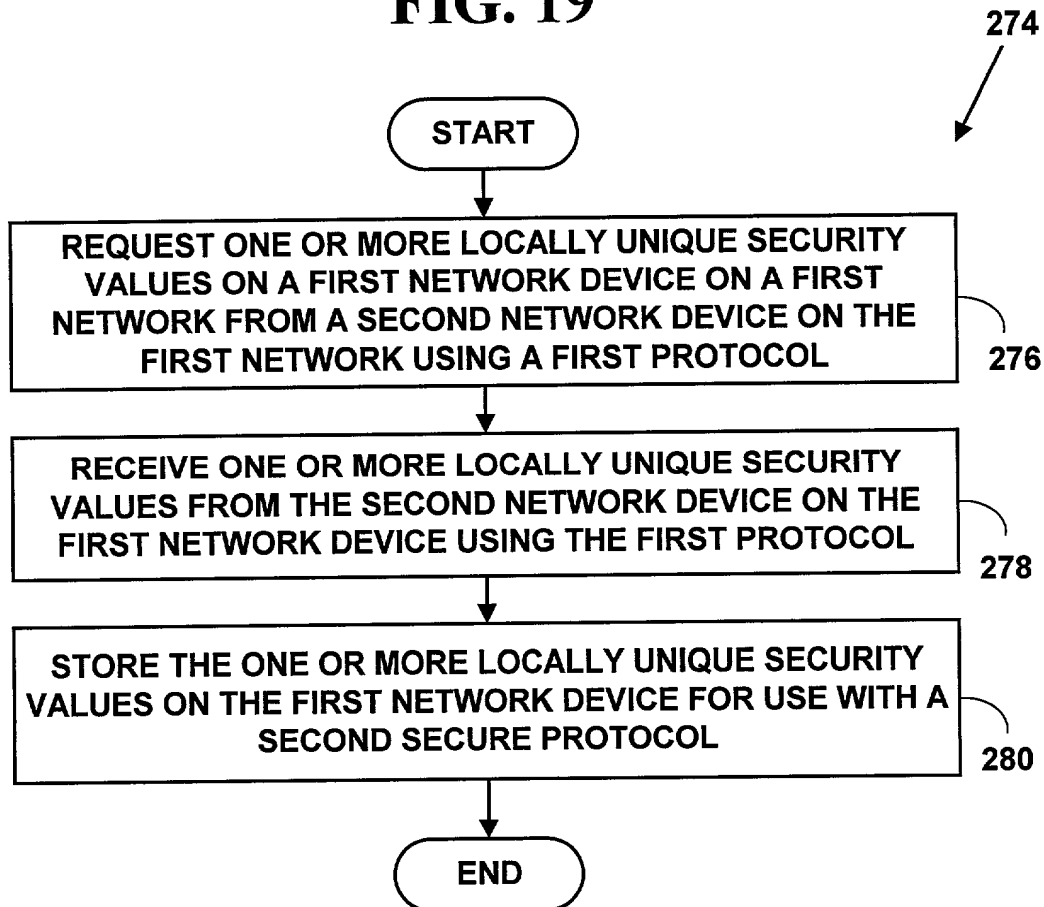


662230-33148E60

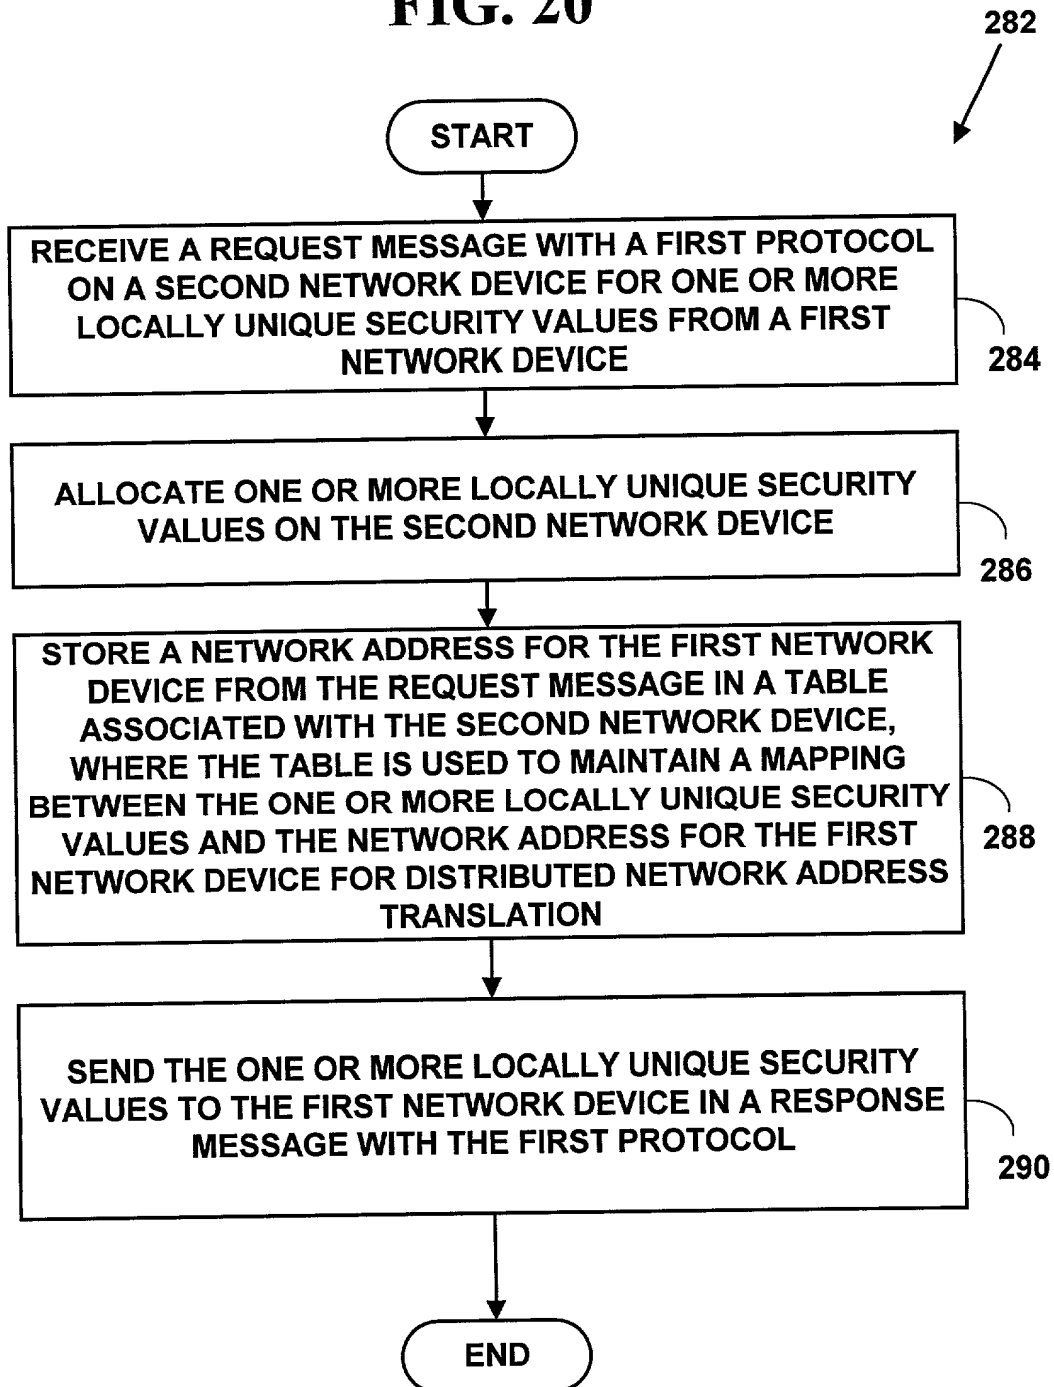
FIG. 18



**FIG. 19**



**FIG. 20**



**FIG. 21**

Diagram illustrating the structure of the SPI-TO-INTERNAL-NETWORK ADDRESS TABLE. The table is divided into three columns: INTERNAL NETWORK ADDRESS (width 294), LOWEST SPI (width 296), and NUMBER OF SPIs (width 298). The table contains two rows of data. The first row shows an internal network address of 10.0.0.1, a lowest SPI of 280, and 32 SPIs. The second row shows an internal network address of 10.0.0.3, a lowest SPI of 312, and 16 SPIs. A pointer 292 indicates the right edge of the table. Vertical dimensions 300 and 302 are shown on the right side of the table rows.

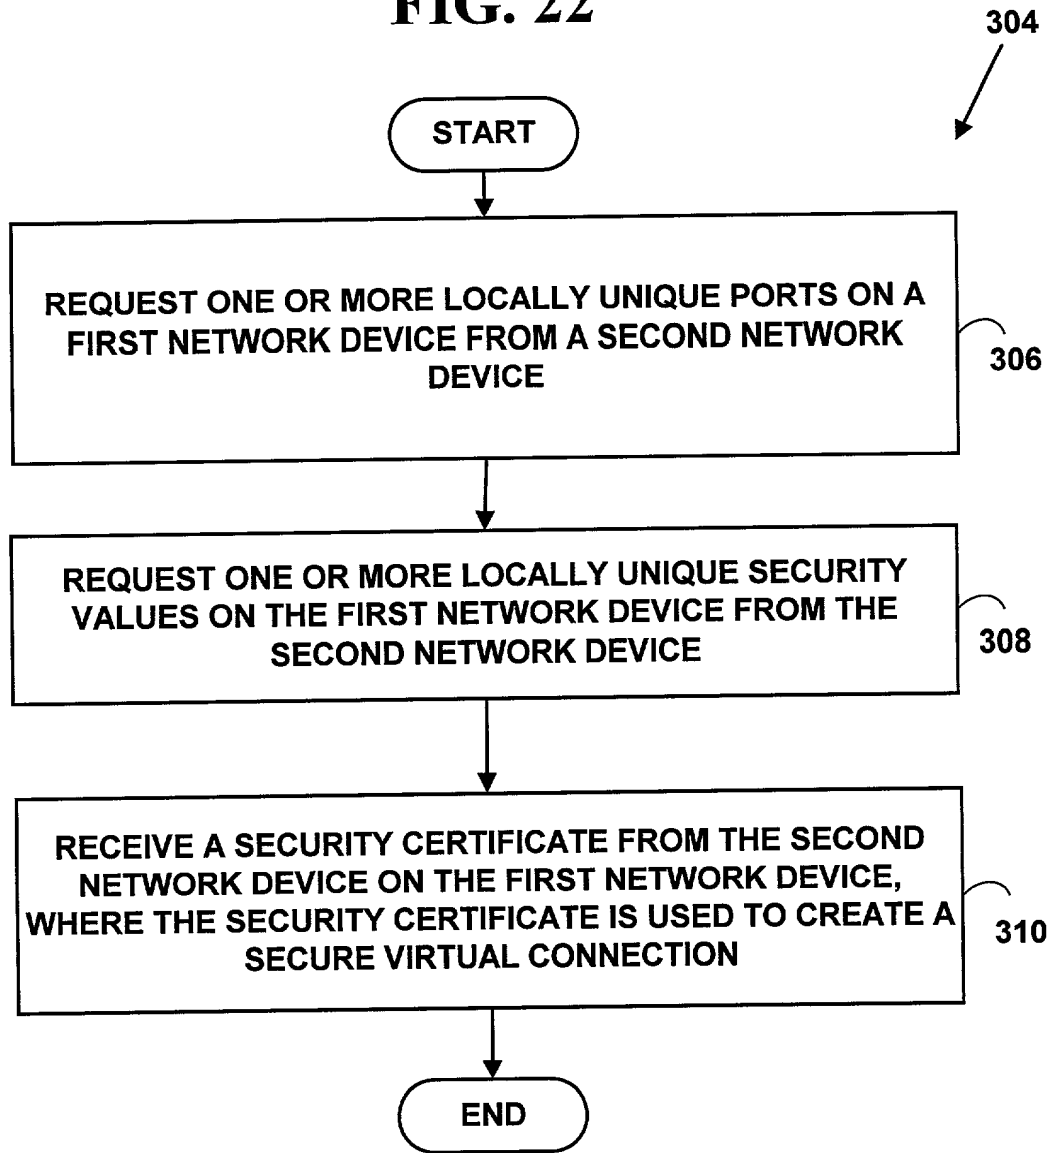
INTERNAL NETWORK ADDRESS	LOWEST SPI	NUMBER OF SPIs
10.0.0.1	280	32
10.0.0.3	312	16

**SPI-TO-INTERNAL-NETWORK ADDRESS  
TABLE**

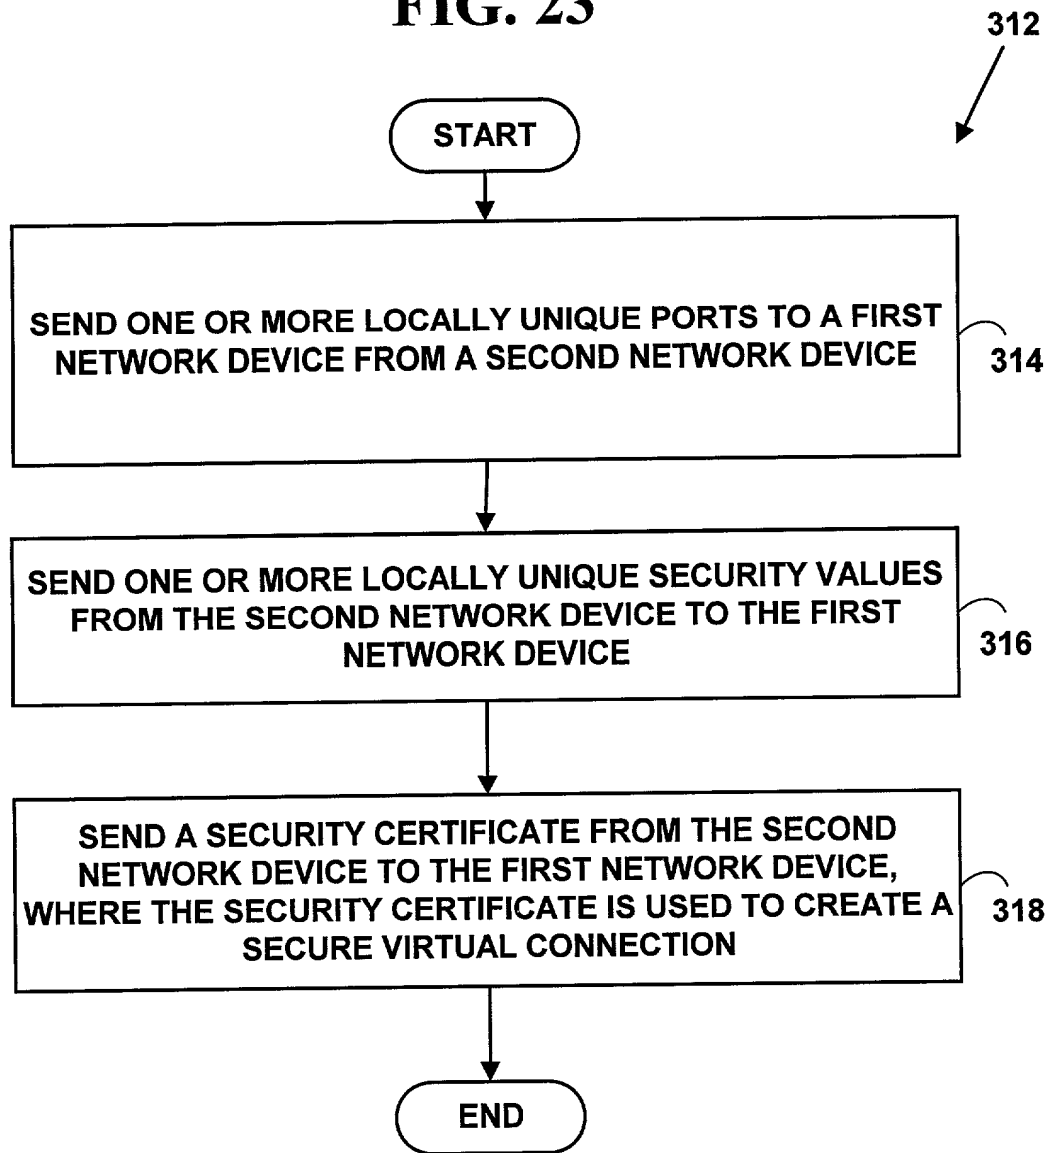
602e30-314860



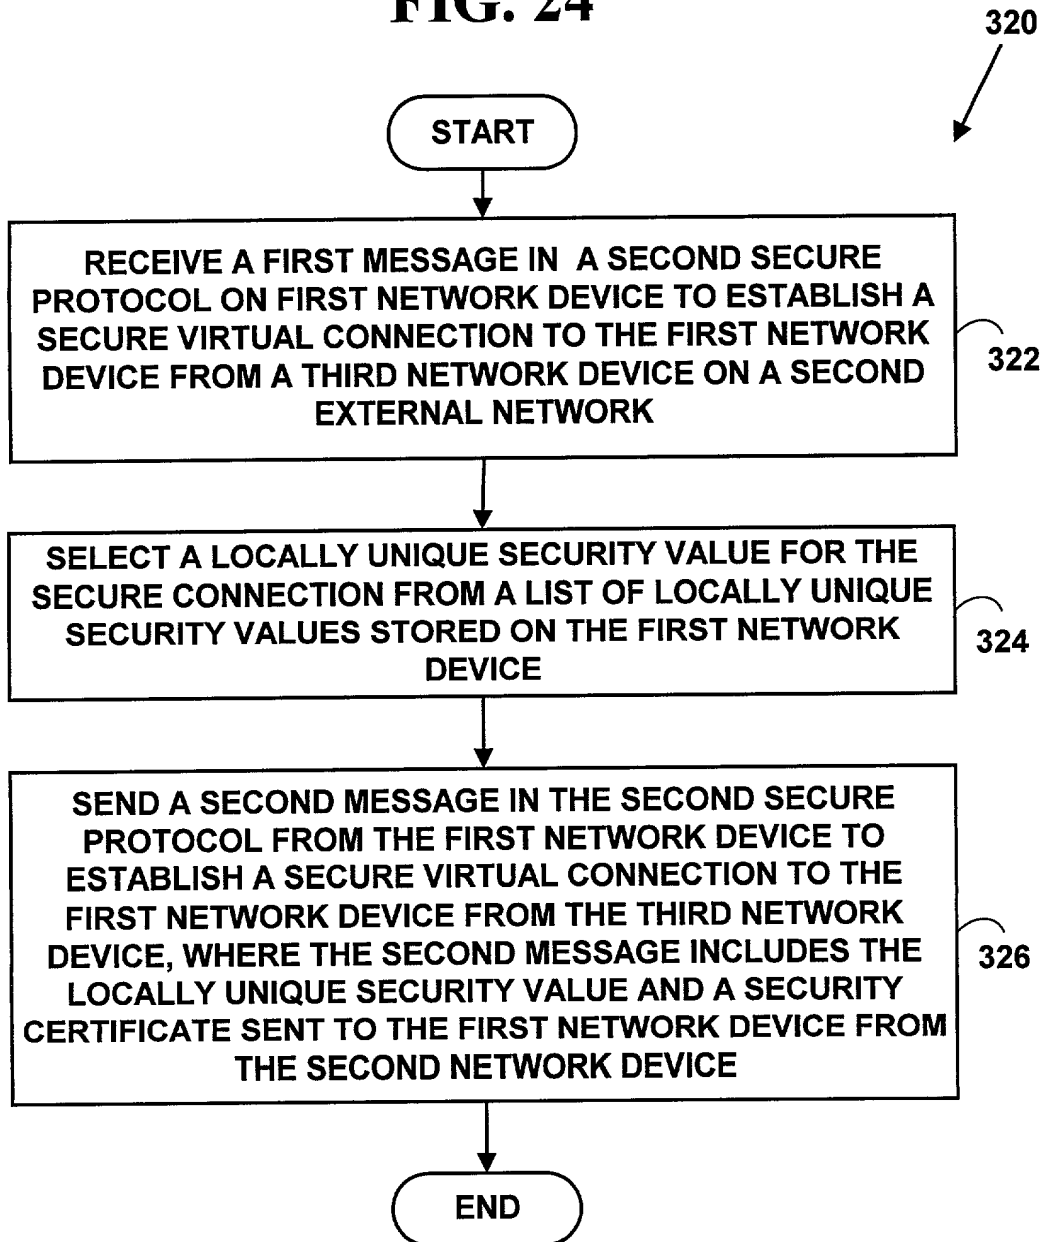
**FIG. 22**



**FIG. 23**



**FIG. 24**



**FIG. 25**

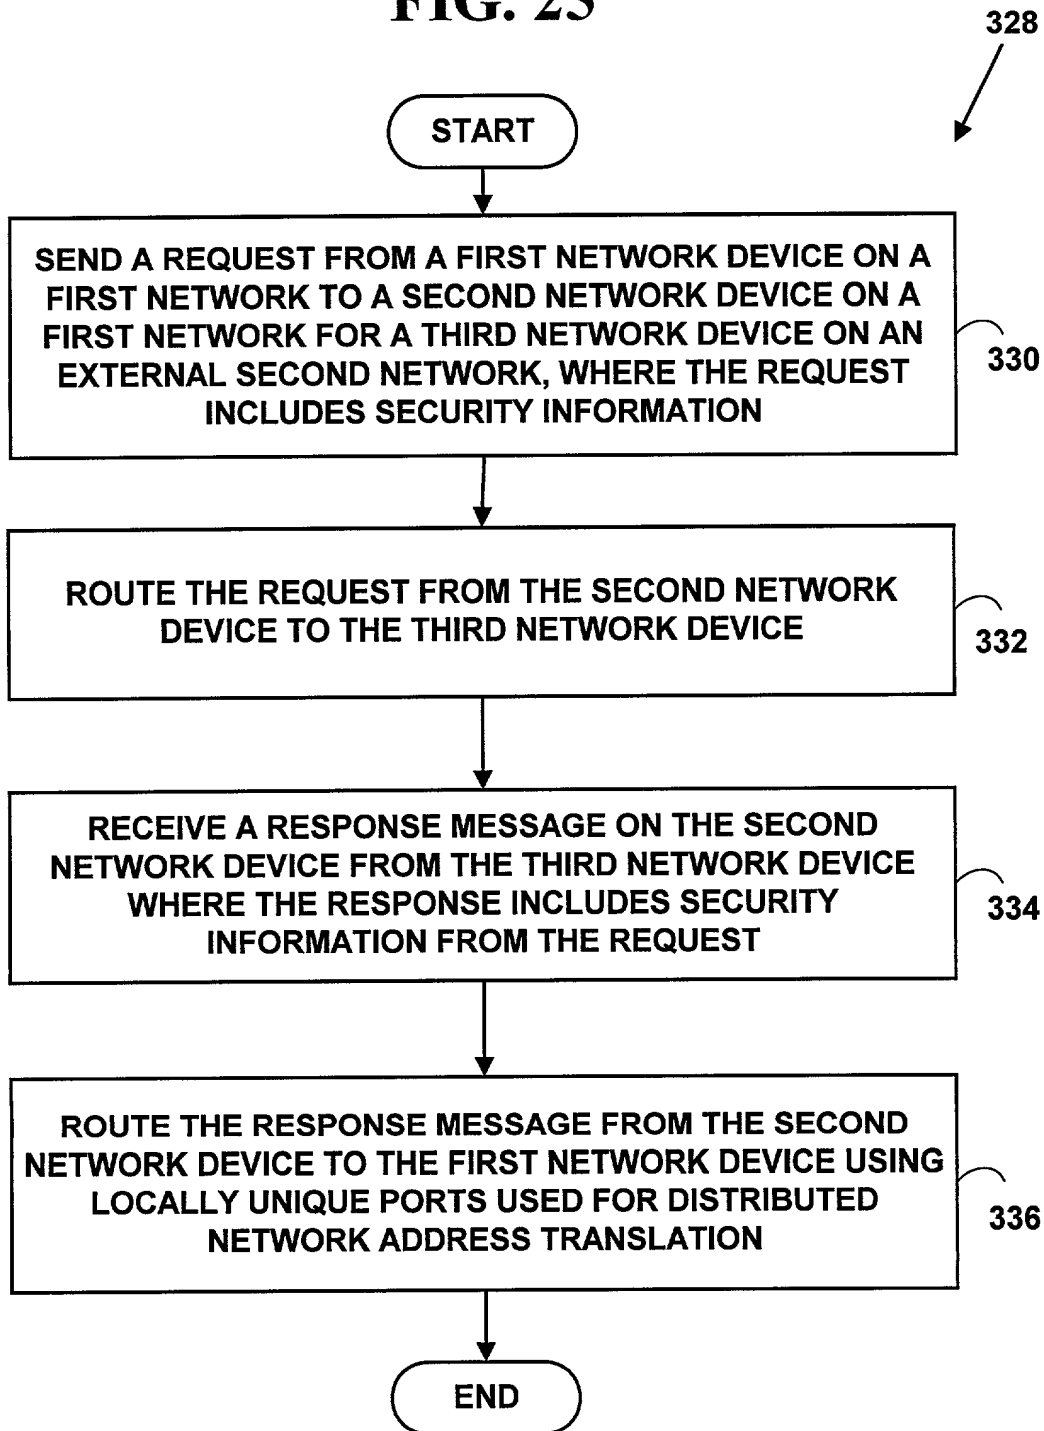

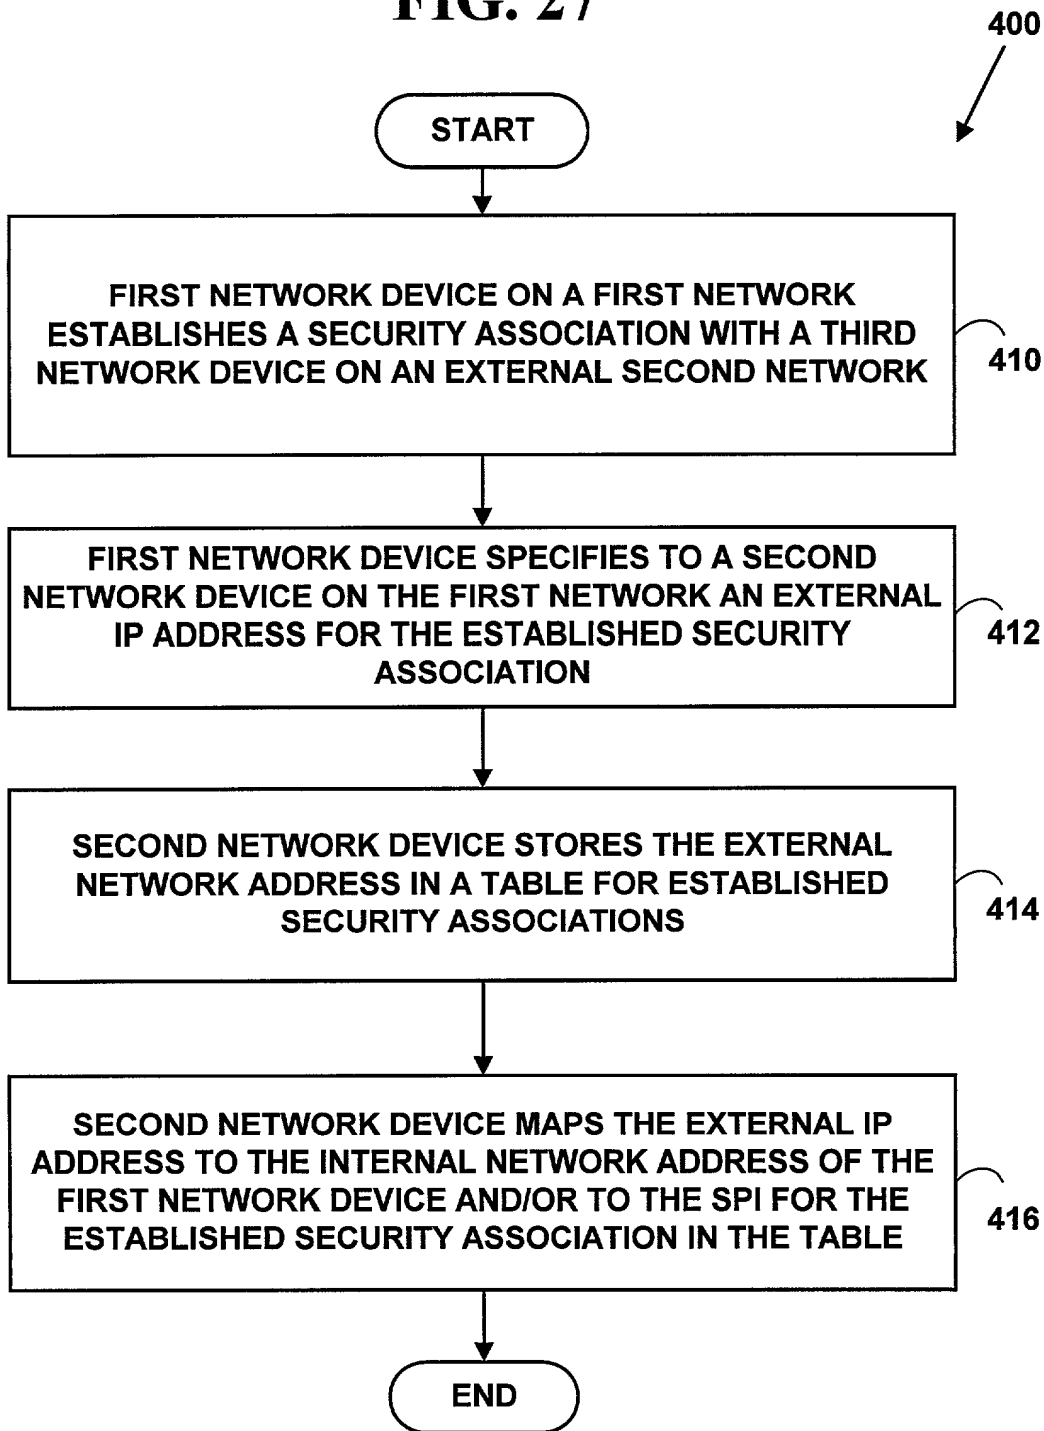



FIG. 26




**FIG. 27**




**FIG. 28A**  
PAP EXTERNAL ADDRESS VALIDATING  
MESSAGE LAYOUT




**FIG. 28B**  
PAP EXTERNAL ADDRESS INVALIDATING  
MESSAGE LAYOUT




**FIG. 29A**



**FIG. 29B**



**FIG. 29C**





**FIG. 30**

